Polarimetric Enhancements to Electro-Optical Aided Navigation Techniques

THESIS

Jeremiah Johnson, Captain, USAF

AFIT/GE/ENG/11-19

AFIT/GE/ENG/11-19

Polarimetric Enhancements to Electro-Optical Aided Navigation Techniques

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science in Electrical Engineering

Jeremiah Johnson, B.S.E.E.

Captain, USAF

March 2011

POLARIMETRIC ENHANCEMENTS TO ELECTRO-OPTICAL AIDED NAVIGATION TECHNIQUES

Jeremiah Johnson, B.S.E.E.

Captain, USAF

Approved:

(signature)	3 MAR 11
Lt Col M.J. Veth, PhD (Chairman)	date
(signature)	7 Mar 2011
Maj K.A. Fisher, PhD (Member)	date
(signature)	7 Mar 11
Dr. M.A. Marciniak (Member)	date

AFIT/GE/ENG/11-19

Abstract

Navigation in indoor and urban environments by small unmanned systems is a topic of interest for the Air Force. The Advanced Navigation Technology Center at the Air Force Institute of Technology is continually looking for novel approaches to navigation in GPS deprived environments. Inertial sensors have been coupled with image aided concepts, such as feature tracking, with good results. However, feature density in areas with large, flat, smooth surfaces tends to be low.

Polarimetric sensors have been used for surface reconstruction, surface characterization and outdoor navigation. This thesis combines aspects of some of these algorithms along with a realistic, micro-facet polarimetric model and a Kalman filter approach to determine surface structure and platform orientation in an indoor environment.

An iterative approach was taken to reach this goal. Several MATLAB graphical user interfaces were developed to determine the ability to estimate surface material parameters. The results of these tests demonstrated the need to constrain the geometry to a specular region. A more complex simulation software package was used to estimate surface orientation given the full set of surface material parameters. An additional set of simplifying assumptions was also developed to reduce the amount of required information. Finally, a physical polarimeter was designed and built to test the algorithms in a realistic environment.

There are three main points that can be taken from this thesis. First, a full set of material parameters can only be determined for a single view by using a multiple hypothesis testing method and only under known geometry conditions. Next, a measurement model for the estimation of pitch angle showed an uncertainty in estimation of $6°$ and a mean error dependent on the material and geometry of a particular situation. Finally, an improvement in attitude estimation of up to 50% was demonstrated.

Acknowledgements

<div align="center">

Table of Contents

</div>

List of Figures

List of Symbols

List of Abbreviations

Polarimetric Enhancements to Electro-Optical Aided Navigation Techniques

I. Introduction

Navigation in indoor and urban environments by small, unmanned systems is a topic of interest for the Air Force. The Advanced Navigation Technology Center at the Air Force Institute of Technology is continually looking for novel approaches to navigation in GPS-deprived environments. GPS-deprived areas require alternate methods of periodic positioning in order to constrain drift in inertial navigation sensors. Inertial sensors have been coupled with image aided concepts, such as feature tracking with good results [7, 27]. However, in areas with large, flat, smooth surfaces, there may not be enough features between frames to make an accurate position update.

Polarimetric sensors have been used for surface reconstruction [10, 12, 37, 38], surface characterization [4, 17] and outdoor navigation [9]. However, these algorithms tend to focus on a single aspect of the authors' respective research areas and require information about the aspect on which they are not focusing. For example, [17] requires knowledge of surface geometry to determine surface characteristics and [38] uses a simplified polarimetric model with known material parameters in order to determine surface structure.

This thesis will combine aspects of some of these algorithms along with a realistic polarimetric model and a Kalman filter approach to determine surface structure and platform orientation in an indoor environment. This chapter continues by presenting the approach to the study and some assumptions made along the way (Section 1.1). It

then presents the set of contributions achieved from this research effort (Section 1.2). Finally, the organization of the rest of the thesis is presented in Section 1.3.

1.1 Approach and Assumptions

An iterative approach was taken in this thesis. First, the estimation of particular parameters of the Shell target polarimetric model, described in Chapter II, were tested individually. The results of these tests showed which parameters were most important in intensity, degree of polarization and angle of polarization measurements and which could be reasonably neglected. Once observable material characteristics were determined, multiple parameter estimation techniques were explored to determine interdependencies in parameter and geometry estimation. These tests show that the full set of polarimetric model parameters and surface geometry can not be determined simultaneously, as expected. However, with a reasonable geometry estimate, a multiple hypothesis testing algorithm proved that surface parameters can be determined given a limited number of materials in a database and a specular geometry.

Given the results of the parameter estimation tests, constraints were placed on subsequent testing geometries, such that only geometries in which a source is present in the specular direction of the camera will be used. This assumption is necessary because any geometry with an off-specular reflection will result in a low degree-of-polarization measurement and will not produce useful measurements.

In order to determine surface orientation with the least amount of knowledge available, a simplification of the Shell target model was created. This simplification makes the assumption that any materials of interest are smooth, dielectric materials. This allows the set of Shell parameters to be narrowed to only the complex index of refraction. Generally, for a smooth dielectric surface, the index of refraction can be reasonably assumed [17].

Finally, in order to determine camera orientation, the Manhattan World constraint was used. For an indoor hallway environment, it was assumed that most large flat surfaces tend to have orthogonal surface normals. Using a Kalman filter approach,

the camera orientation is known with limited uncertainty. This information is used to determine surface parameters by way of the multiple hypothesis testing method, and then used to determine relative geometry with surfaces in the scene. This relative geometry is then used to update errors in the filter, and the cycle is repeated.

Three sets of tools are used throughout the thesis. These tools are described in detail in Chapter III. The same iterative approach was taken using these tools. In general, MATLAB graphical user interfaces are used to test a hypothesis. Then, DIRSIG simulation software is used to test the algorithm in a more complex environment. Finally, a physical polarimeter is used to determine if any anomalies exist under real world conditions.

1.2   Contributions

The contributions from this research effort apply broadly to both polarimetric and navigation technology fields of study. Because of the continued research efforts in polarimetry, as mentioned above, the contributions from the tests to determine surface parameters can be useful in image cuing and target detection algorithms. However, the main contribution focus is the addition of a polarimetric measurement model to a Kalman filter algorithm, common in navigation.

This modification is powerful in its simplicity and availability. The addition of the measurement model to existing Kalman filters is a simple software update. Several methods have been used to add a polarimetric capability to an existing camera, and these additions add little or no weight or power requirements to the existing systems. The polarimetric imagery complements feature tracking EO-aiding algorithms well, because it thrives in environments in which feature tracking algorithms do not. Because each view is independent of previous views, it can be used to constrain drift in an inertial system. Finally, by having an estimate of the camera orientation, previously established algorithms can be used to determine additional surface structure.

1.3 Organization

This thesis is organized as follows. Technical background information is presented in Chapter II. A description of the simulation software, design of a set of MATLAB graphical user interfaces, and design and construction details of a physical polarimeter are presented in Chapter III. Test methodology and set up for each experiment is presented in Chapter IV. Results from these tests are analyzed in Chapter V. Finally, a set of conclusions and suggestions for future work is presented in Chapter VI.

II. Technical Background

This chapter presents technical background information required to understand the material posed throughout the research. Variable notation is detailed in Section 2.1. Section 2.2 will discuss various frames of reference common to navigation and how to transform from one frame to another. Current electro-optically aided navigation techniques, used to constrain drift in inertial navigation systems, is presented in Section 2.3. A brief background on polarimetry is given in Section 2.4, with discussion of a few polarimetric models in Section 2.5. Section 2.6 will discuss some polarimetric shape recovery techniques already in use. Finally, a brief overview of the Unscented Kalman Filter is presented in Section 2.7.

2.1 Variable Notation

This section describes the variable types and notation used throughout this thesis.

- **Scalars** are represented by upper or lower case letters in italic type (e.g., x or X).

- **Vectors** are represented by lower case letters with bold font, (e.g., $\boldsymbol{x}$). Each vector is composed of a column of scalar elements denoted by x_i, where i represents the element number.

- **Homogeneous Vectors**, vectors in which the last element is a 1, are denoted by an underline (e.g., $\underline{x}$).

- **Matrices** are given as upper case letters in bold font. The matrix $\boldsymbol{X}$ is composed of elements X_{ij} where i is the row index and j is the column index.

- **Direction Cosine Matrices** from frame a to frame b are given as $\boldsymbol{C}_a^b$.

- **Reference Frames** are described by superscripts. For example, $\boldsymbol{p}^a$ is a vector expressed in the a frame.

- **Mean** values are defined by a bar, such as $\bar{x}$.

- **Covariance** of variables is defined by a capital $\boldsymbol{P}$ and two subscripts. The subscripts describe the variables with which the covariance is taken. (e.g., $\boldsymbol{P}_{xy} = \sum(x_i - \bar{x})(y_i - \bar{y})$).

2.2 Frames of Reference

In navigation, relative positions of two objects need to be expressed in a common coordinate system. For example, the location of an object in an image can be expressed in the camera's frame of reference. The location and orientation of the camera relative to a local navigation frame can then be used to help determine a line along which the image feature is located in a localized coordinate system. In this section, a few common coordinate systems will be presented along with a technique for converting vectors from one system to another.

2.2.1 Coordinate Systems. A coordinate system can be defined at any location and orientation. For this research, most coordinates will be expressed in terms of a camera location and orientation in the local navigation frame. The local navigation frame is generally described by a single position on the Earth and an orientation. The orientation of the local navigation frame axes is arbitrary and can be defined for a given problem.

Features and surfaces within an image are described in the camera coordinate system. Figure 2.1 shows a typical camera coordinate system. This system is described as being located at the optical center of the camera with the x-axis parallel to the focal plane and pointed up, the y-axis parallel to the focal plane and pointed out the right hand side and the z-axis perpendicular to the focal plane and pointed out the center of the lens of the camera.

Using a camera coordinate system allows for a simple way to locate a feature in an image. However, in order to build a 3-D model of a scene, feature locations in the images must be transformed into a common coordinate system. The next section will discuss how to convert vectors between coordinate systems.
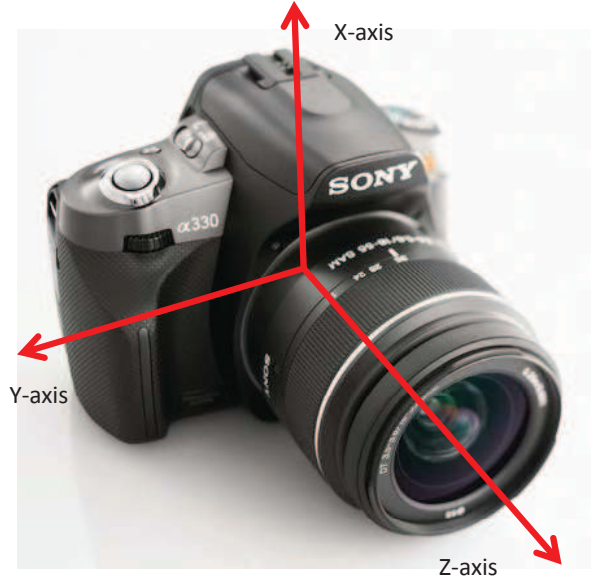
Figure 2.1:   Common coordinate system for a camera reference frame. This figure shows the x-axis pointed out the top of the camera, the y-axis pointed out the right hand side and the z-axis pointed out the front of the camera from the center of focus through the optical center of the lens.

2.2.2   Coordinate Conversions.   There are several methods of converting position and orientation vectors from one frame of reference to another. Two types will be presented in this section, Direction Cosine Matrix (DCM) Transformations and Euler Angles.

2.2.2.1   Direction Cosine Matrix.   The DCM is a 3×3 matrix whose columns represent vectors in the body axis projected onto a reference axis [32]. The DCM presents a convenient way to transform vectors in one coordinate system into another by simple multiplication. A vector quantity defined in body axes may be expressed in another reference axes by pre-multiplying the vector by the DCM. Equation (2.1) shows how this is done, where $\boldsymbol{p}^a$ is the position of an object in the 'a' reference frame, $\boldsymbol{p}^b$ is the position of the object in the 'b' frame, and $\boldsymbol{C}_b^a$ is the DCM to convert from reference frame 'b' to frame 'a'.

$$\boldsymbol{p}^a = \boldsymbol{C}_b^a \boldsymbol{p}^b \tag{2.1}$$

7

Some useful properties of the DCM include [27]:

1. $Det(\boldsymbol{C}_a^b) = 1$

2. $\boldsymbol{C}_b^a = (\boldsymbol{C}_a^b)^{-1} = (\boldsymbol{C}_a^b)^T$

3. $\boldsymbol{C}_a^c = \boldsymbol{C}_b^c \boldsymbol{C}_a^b$

A DCM may be computed using individual rotations about orthogonal axes in the body frame. These rotation angles are known as Euler angles and are covered in the following section.

2.2.2.2   Euler Angles.   Euler angles are a set of transformations which relate to single rotations around each orthogonal axis in turn [32]. These transformations are then multiplied together to obtain the full transformation. They are commonly used when converting from a body frame to a local navigation frame by using the roll (θ), pitch (ϕ) and yaw (ψ) angles of the body. The three rotations may be expressed as three separate DCMs presented in Equations (2.2) - (2.4).

$$\boldsymbol{C}_n^1 = \begin{bmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2.2}$$

$$\boldsymbol{C}_1^2 = \begin{bmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{bmatrix} \tag{2.3}$$

$$\boldsymbol{C}_2^b = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & \sin\phi \\ 0 & -\sin\phi & \cos\phi \end{bmatrix} \tag{2.4}$$

The DCM from the reference to the body axis can then be expressed as the product of the individual transformations.

$$C_n^b = C_2^b C_1^2 C_n^1 \qquad (2.5)$$

Once a common coordinate system is set, sensors must be used to determine motion through the reference frame. The next section will discuss electro-optically aided navigation techniques, used to determine location and orientation of an vehicle in motion.

2.3 Electro-Optically Aided Navigation

Because of the drift errors associated with inertial navigation systems, described in detail in [32], electro-optical (EO) aiding algorithms have been developed to assist in navigation. These types of algorithms are meant to constrain the errors in inertial sensors by providing periodic position information. Several algorithms exist for EO-aided navigation [7, 19, 27]. The focus of this section is to describe the basics of a feature matching algorithm, which allows for relative positioning between frames by keeping track of a set of invariant features in the scene.

Two sets of geometric constraints used to determine camera motion are discussed. The epipolar constraint is covered in Section 2.3.1, and the homographic constraint is discussed in Section 2.3.2. Each of these constraints requires point correspondence between images. A feature detection and correspondence algorithm is presented in Section 2.3.3. Finally, because these constraints require a pin-hole camera model, a camera calibration technique is presented in Section 2.3.4.

2.3.1 Epipolar Constraints . The epipolar geometry between two views is composed of the geometry of the intersection of the image planes with the pencil of planes having the baseline as axis [16]. The baseline is the line joining the camera centers. The epipole is the point where the vector from one camera to the other camera intersects with the image plane. Figure 2.2 illustrates the epipolar geometry relationships.

9
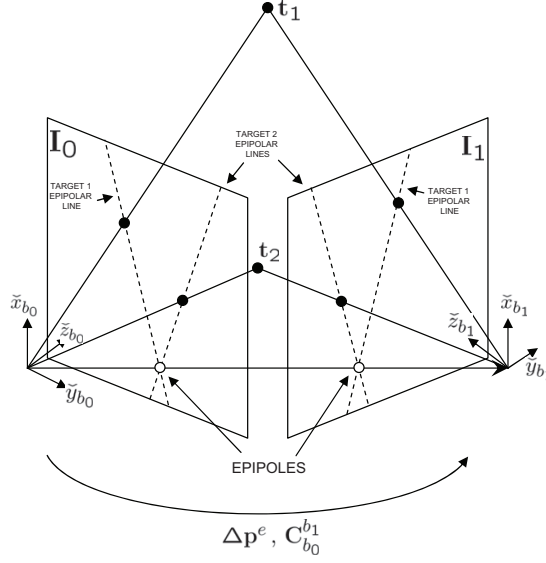
Figure 2.2:   Two view geometry example. This figure shows the relationship between a 3D point in space and two cameras. The epipole is shown as the intersection of the line which connect the focal points of each image with each image plane. [27]

This can be used to better estimate matching features between images. A point in the second image must lie on the plane constrained by the the baseline and the vector pointing from the center of the camera to the image feature. The line in the second image which is produced by the intersection of this plane and the focal plane of the second camera is known as the epipolar line. A matching feature in the second image must lie on this line.

2.3.1.1   Fundamental Matrix .   By looking at the triangle made up of the two camera locations and an object in each scene, it can be seen that there is a simple relationship between the vector from the first camera to the object, $\boldsymbol{s}^a$, and the vector from the second camera to the object, $\boldsymbol{s}^b$. This relationship is shown in Equation (2.6).

$$\boldsymbol{s}^b = \boldsymbol{p}_{ba}^b + \boldsymbol{C}_a^b \boldsymbol{s}^a \tag{2.6}$$

10

In this equation, $\boldsymbol{C}_a^b$ is a DCM from camera a to camera b. The translation vector from camera b to camera a, $\boldsymbol{p}_{ba}^b$, is the vector pointing from camera b to camera a, referenced to the camera b coordinate system.

These vectors are all illustrated in Figure 2.2. The vectors $\underline{\boldsymbol{s}}^a$ and $\underline{\boldsymbol{s}}^b$ are then defined to be the homogeneous $\boldsymbol{s}$ vectors, or the $\boldsymbol{s}$ vector scaled so that $\underline{s}_3 = 1$. Their relationship is therefore, $\boldsymbol{s}^a = \lambda_a \underline{\boldsymbol{s}}^a$, where λ_a is a scaling parameter. Using this relationship, Equation (2.6) can then be rewritten as:

$$\lambda_b \underline{\boldsymbol{s}}^b = \boldsymbol{p}_{ba}^b + \boldsymbol{C}_a^b \lambda_a \underline{\boldsymbol{s}}^a \tag{2.7}$$

Pre-multiplying by the cross product of $\boldsymbol{p}_{ba}^b$ and the dot product of $\underline{\boldsymbol{s}}_b$ will yield Equation (2.8).

$$(\underline{\boldsymbol{s}}^b)^T (\boldsymbol{p}_{ba}^b \times) \boldsymbol{C}_a^b \underline{\boldsymbol{s}}^a = 0 \tag{2.8}$$

The middle term of Equation (2.8), $(\boldsymbol{p}_{ba}^b \times) \boldsymbol{C}_a^b$, is known as the Fundamental matrix. These terms can also be written in the more useful form, $\boldsymbol{C}_b^a (\boldsymbol{p}_{ab}^a \times)$, where $\boldsymbol{C}_b^a$ is the DCM from camera b to camera a and $\boldsymbol{p}_{ab}^a$ is the translation of the camera from position a to position b represented in the camera a coordinate system.

2.3.1.2 Determining the Fundamental Matrix .   It has been shown that the Fundamental matrix is constrained by

$$(\underline{\boldsymbol{s}}^b)^T \boldsymbol{F} \underline{\boldsymbol{s}}^a = 0 \tag{2.9}$$

for any pair of matching points $\underline{\boldsymbol{s}}^a$ and $\underline{\boldsymbol{s}}^b$ [16]. Given enough matches, this equation can be used to compute the unknown matrix $\boldsymbol{F}$. By writing $\underline{\boldsymbol{s}}^a = (x, y, 1)^T$ and $\underline{\boldsymbol{s}}^b = (x', y', 1)^T$, each match gives a linear equation in the unknown entries of $\boldsymbol{F}$. The coefficients of this equation are easily written in terms of the known coordinates $\underline{\boldsymbol{s}}^a$ and $\underline{\boldsymbol{s}}^b$. Specifically, the equation corresponding to the pair of points is:

$$x'xf_{11} + x'yf_{12} + x'f_{13} + y'xf_{21} + y'yf_{22} + y'f_{23} + xf_{31} + yf_{32} + f_{33} = 0 \quad (2.10)$$

Using this relationship, the algorithm for determining the Fundamental matrix from a set of eight or more matching points can be broken down into four steps. The first step requires normalization of the image coordinates. Then, a linear solution to the matrix, $\hat{\boldsymbol{F}}'$, is determined by finding the singular vector corresponding to the smallest singular value of $\boldsymbol{A}$, where $\boldsymbol{A}$ is determined by Equation (2.11).

$$\boldsymbol{A}f = \begin{bmatrix} x'_1x_1 & x'_1y_1 & x'_1 & y'_1x_1 & y'_1y_1 & y'_1 & x_1 & y_1 & 1 \\ x'_2x_2 & x'_2y_2 & x'_2 & y'_2x_2 & y'_2y_2 & y'_2 & x_2 & y_2 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x'_nx_n & x'_ny_n & x'_n & y'_nx_n & y'_ny_n & y'_n & x_n & y_n & 1 \end{bmatrix} f = 0 \quad (2.11)$$

A constraint is then placed on the matrix such that $|\hat{\boldsymbol{F}}'| = 0$. This is done by performing a singular value decomposition and recreating the matrix using only the two largest singular values resulting in $\hat{\boldsymbol{F}}$. Finally, a denormalization is performed using the normalizing transformations found in the first step which yields the final Fundamental matrix, $\boldsymbol{F}$.

This Fundamental matrix algorithm is not very robust if it simply uses all 'matches' between images. A small mismatch will throw off the final DCM and position vector. The Random Sample Consensus, RANSAC, algorithm can be used to further refine the matches found in Section 2.3.3. This algorithm starts by selecting a random sample of normalized matching pairs and finds a Fundamental Matrix for this set. If it finds an acceptable Fundamental Matrix for the set, it will then add and evaluate the rest of the samples against the proposed matrix. If an acceptable number of samples correspond with the proposed model, the RANSAC function will return

the best fit for the function and the pairs that match the function within a threshold. Otherwise, the function will start the process over with a new set of initial samples. This continues until all options have been exhausted or an acceptable solution is found.

2.3.1.3   Decomposition of the Fundamental Matrix .   The decomposition of the Fundamental Matrix is what allows for the determination of camera motion. Once the Fundamental Matrix is determined, a singular value decomposition can be performed to decompose the 3×3 Fundamental Matrix into a 3×3 DCM and a 3×1 translation vector. A singular value decomposition converts the Fundamental Matrix into a diagonal matrix, $\boldsymbol{S}$, and unitary matrices, $\boldsymbol{U}$ and $\boldsymbol{V}$. The matrix, $\boldsymbol{W}$, a DCM for a $\pi/2$ rotation about the z-axis, is also required. Equation (2.12) shows how these can be used to formulate the DCM, $\boldsymbol{C}_b^a$. Equation (2.13) shows one way to produce the translation vector.

$$\boldsymbol{C}_b^a = \pm\boldsymbol{U}\boldsymbol{W}\boldsymbol{V}' or \pm \boldsymbol{U}\boldsymbol{W}'\boldsymbol{V}' \tag{2.12}$$

$$\boldsymbol{p}_{ab}^a = \pm\boldsymbol{U}\boldsymbol{Z}\boldsymbol{U}' \tag{2.13}$$

where $\boldsymbol{Z}$ is defined to be:

$$\boldsymbol{Z} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{2.14}$$

These calculations yield four combinations of DCM and position vector. The correct pair is the pair that produces only positive λ values in Equation (2.7). These λs can be calculated by manipulating Equation (2.7) into the form of Equation (2.15).

$$- \boldsymbol{C}_a^b \boldsymbol{p}_{ab}^a = \left[\begin{array}{cc} \boldsymbol{C}_a^b \underline{\boldsymbol{s}}^a & -\underline{\boldsymbol{s}}^b \end{array} \right] \left[\begin{array}{c} \lambda_a \\ \lambda_b \end{array} \right] \tag{2.15}$$

Equation (2.16) uses a least squares technique to find the estimate for λ, in which $\boldsymbol{L} = \left[\begin{array}{cc} \boldsymbol{C}_a^b \underline{\boldsymbol{s}}_a & -\underline{\boldsymbol{s}}_b \end{array} \right]$

$$\left[\begin{array}{c} \hat{\lambda}_a \\ \hat{\lambda}_b \end{array} \right] = (\boldsymbol{L}^T \boldsymbol{L})^{-1} \boldsymbol{L}^T \left[\begin{array}{c} \underline{\boldsymbol{s}}_a \\ \underline{\boldsymbol{s}}_b \end{array} \right] \tag{2.16}$$

For a set of coplanar features, a special geometry constraint exists, known as a homography, which can be used to more easily correspond points between images. Section 2.3.2 describes how the homographic matrix is calculated, and how it can be used to assist in determining camera rotation and translation parameters or a flat surface normal.

2.3.2   Homographic Geometry .   Through the homography a point in one view determines a point in the other which is the intersection of the pointing vector with the plane [16]. Given a set of matching points on a plane, the homographic transformation for a pin-hole model is given as

$$\boldsymbol{s}^b = \boldsymbol{H} \boldsymbol{s}^a \tag{2.17}$$

The Homographic transformation matrix, $\boldsymbol{H}$, can be found by using a method similar to that used for the Fundamental matrix in which a RANSAC algorithm is fed a set of matching points and a best fit to the above equation is produced. In order to use the following equations, this $\boldsymbol{H}$ matrix must be normalized. Equation (2.18) shows how to normalize the $\boldsymbol{H}$ matrix by taking the singular value decomposition, where σ_2 is the second singular value of $\boldsymbol{H}$.

$$\boldsymbol{H}_{norm} = \frac{\boldsymbol{H}}{\sigma_2} \tag{2.18}$$

14
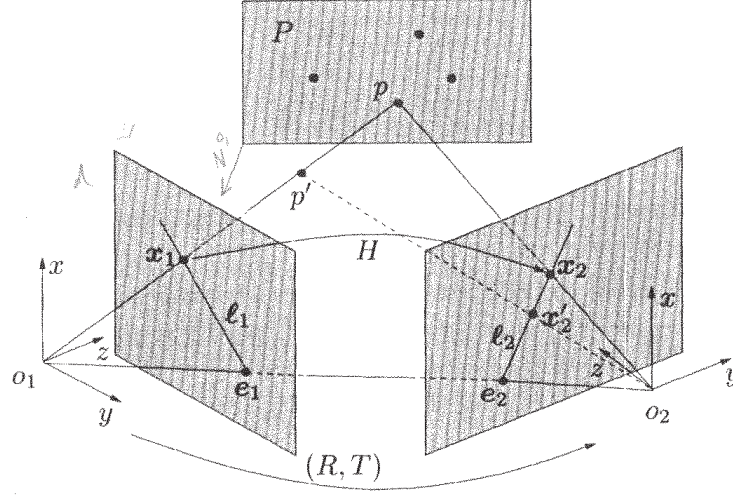
Figure 2.3:   Homographic Geometry Illustration. This figure shows how the series of epipolar planes can be uniquely defined for a plane in the image [27].

Since Equation (2.6) still holds for any single point in the scene, it can be rewritten as

$$\boldsymbol{s}^b = \boldsymbol{C}_a^b \boldsymbol{s}^a + \boldsymbol{p}_{ba}^b \tag{2.19}$$

in which $\boldsymbol{C}_a^b$ refers to the rotation matrix from the camera a frame to the camera b frame. $\boldsymbol{p}_{ba}^b$ is the vector describing the translation from the a frame to the b frame, represented in the b frame. Figure 2.3 shows an example of the the homographic relationship.

It can then be shown that there exists a relationship between the homography matrix, $\boldsymbol{H}$, the rotation and translation of the camera, and the surface normal vector, $\boldsymbol{n}$. This relationship is expressed in Equation (2.20).

$$\boldsymbol{H}_{norm} \dot{=} \boldsymbol{C}_a^b + \frac{1}{d}\boldsymbol{p}_{ab}^a \boldsymbol{n}^T \tag{2.20}$$

2.3.3   Correspondence Between Images .   Object correspondence between images is a requirement for both sets of geometric constraints presented in the previ-

ous sections. Several ways have been proposed for image point correspondence. The method used throughout this thesis begins with the Scale Invariant Feature Transformation (SIFT) algorithm [20,21]. This particular algorithm categorizes image features by a feature descriptor vector.

There are many ways to match these features between images. For the feature descriptor vector given by the SIFT algorithm, a vector dot product is the most convenient. The dot product of two large dimensional vectors like these will show how well the vectors align. The pixel locations of the keypoints with the best alignment, within a threshold, can be considered to be matching features. However a matching algorithm which simply uses the largest dot product within a threshold can lead to a small number of mismatches. In order to filter out mismatches, a minimal difference measurement between the best two matches can be implemented.

Once a set of good matches has been found, the pixel locations still need to be converted to vectors pointing to the features. Section 2.3.4 will show how distortion is removed and pixel locations are converted to normalized pointing vectors.

2.3.4 Camera Calibration . Because line intersections are used to determine feature locations in the local navigation frame and most camera lenses will distort these lines in a radial pattern, a camera calibration must be done to remove these distortions. For the physical polarimeter, described in Chapter III, the Caltech Calibration Kit was used. This calibration kit is described in great detail on their website [6].

The output from this calibration is the set of intrinsic camera parameters:

- Focal length (fc)

- Principal point (cc)

- Skew (αc)

- Distortion (kc)

- Pixel error (err)

Once a calibration is performed, these parameters can now be used to undistort any images that are shot with the same lens parameters used in the calibration. The intrinsic camera matrix presented in Equation (2.21) is a transformation matrix from a line of sight vector to normalized pixel location.

$$\boldsymbol{T}_{los}^{pix} = \begin{bmatrix} fc(1) & \alpha_c fc(1) & cc(1) \\ 0 & fc(2) & cc(2) \\ 0 & 0 & 1 \end{bmatrix} \tag{2.21}$$

The inverse of $\boldsymbol{T}_{los}^{pix}$, T_{pix}^{los}, can be used to transform from normalized pixel location to line of sight vector.

In areas with minimal features, an alternative approach to EO-aiding must be applied. The next section describes the phenomenology of polarimetry and why it is useful in such a situation.

2.4 Polarimetric Phenomenology

Polarization is the measurement of the electric field vector orientation in electromagnetic radiation. The electric field of light, $E(t)$, can be defined by a time varying vector represented in Equation (2.22). E_0 is the amplitude of the electromagnetic field, ω is the angular frequency, $\boldsymbol{k}$ is the wave vector, t is time, $\boldsymbol{z}$ is the direction of propagation and ϕ is a constant phase shift [14].

$$E(t) = E_0 \sin(\boldsymbol{k} \cdot \boldsymbol{z} - \omega t + \phi) \tag{2.22}$$

Figure 2.4 shows an illustration of a light wave. The plane orthogonal to the direction of travel of the light wave will be defined to have an arbitrary 0° point and an angle relative to this point, θ, representing the angle of the oscillation of the electric wave and will be used to describe the orientation of the electric field.

17
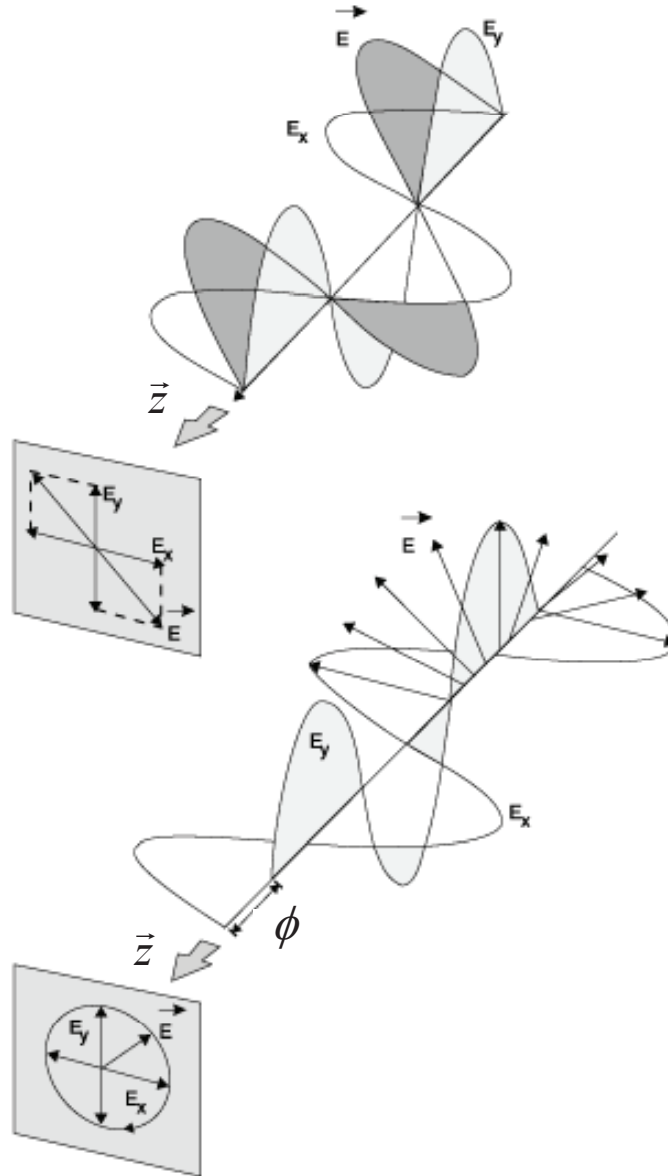
Figure 2.4: Electro-Magnetic Field Representation. The phase offset of the electric and magnetic components of the wave directly relate to the polarization of light in the wave. [15]
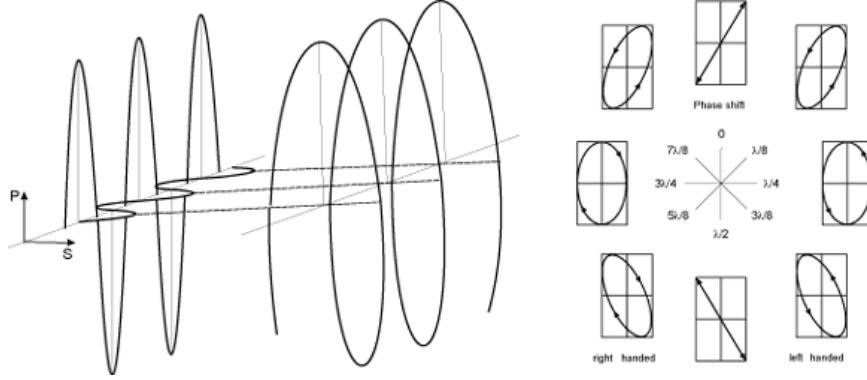
Figure 2.5: Representation of Linear, Circular, and Elliptical Polarization. Projections of the 2-dimensional electromagnetic vector onto the transverse path create different degree and angle of polarization states based on the phase difference between the two orthogonal components. [14]

2.4.1 Polarization States . The polarization state of a coherent light wave represents the direction of oscillation of the electric field in the plane perpendicular to the direction of motion. When the perpendicular components of the electric field oscillate in phase, this is known as linear polarization. When these components oscillate with the same amplitude but at 90° out of phase, this is known as circular polarization. Otherwise, the resulting polarization state is said to have elliptical polarization, in that the shape traced in the x-y plane through a full oscillation cycle is an ellipse [14]. Figure 2.5 shows a representation of these states.

In a coherent light source, the polarization can be easily classified because there is only one set of waves. However, a typical sensor will detect incoherent light in most situations. Incoherent light will consist of a combination of polarization states. In 1852, George Gabriel Stokes developed a system for describing the polarization state of incoherent radiation [14]. Section 2.4.2 covers the background of the Stokes vector.

2.4.2 Stokes Vector . The Stokes vector, $\boldsymbol{S}$, is a 4-element column vector which represents the polarimetric information in incoherent light. Equation (2.23) shows the makeup of the vector.

$$\boldsymbol{S} = \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix} \tag{2.23}$$

The S_0 component represents the total irradiance of light incident on an object, such as a detector. It can be represented as the sum of the electric field in the $0°$ and $90°$ directions of the coordinate system described in Section 2.4.1, or $S_0 = E_0 + E_{90}$. S_1 represents the amount of polarization in the $0°$ or $90°$ direction and is calculated by $S_1 = E_0 - E_{90}$. A positive S_1 represents light that is more polarized in the $0°$ direction, while a negative S_1 describes light that is more polarized in the $90°$ orientation. The S_2 element represents the amount of polarization between $45°$ and $135°$ and is defined to be $S_2 = E_{45} - E_{135}$. Finally, the S_3 element represents the amount of circularly polarized light by, $S_3 = E_{rc} - E_{lc}$, where E_{rc} represents clockwise circular polarization and E_{lc} represents counter-clockwise circular polarization. These elements can then be combined to describe the amount of linearly polarized light (DoLP) calculated using Equation (2.24), and the angle of polarization (AoP), Φ, found in Equation (2.25).

$$DoLP = \frac{\sqrt{S_1^2 + S_2^2}}{S_0} \tag{2.24}$$

$$\Phi = \frac{1}{2} \arctan \frac{S_2}{S_1} \tag{2.25}$$

These equations are the foundation of polarimetry for remote sensing. In order to estimate a set of Stokes Vectors from the reflection of a surface, a model must be constructed. Section 2.5 discusses some polarimetric models used in modern simulation software.

2.5 Polarimetric Models

A method of converting Stokes parameters from just before to just after a reflection is known as Mueller matrix calculus and is discussed in Section 2.5.1. Several models have been proposed to find Mueller matrices for a wide variety of scenarios. Each model uses the set of Fresnel equations described in Section 2.5.2. This basic model is improved upon with the addition of a Bidirectional Reflectance Distribution Function, BRDF, and micro-facet models presented in Sections 2.5.3 and 2.5.4. The full model used in this paper is known as the Shell Target model and is described in detail in Section 2.5.5.

2.5.1 Mueller Matrix . A Mueller matrix, $\boldsymbol{M}$, is a 4×4 matrix describing the transition of one state of Stokes vectors to another at an interface. The Mueller matrix conversion is found in Equation (2.26), in which $\boldsymbol{S}_i$ refers to the incident Stokes vector and $\boldsymbol{S}_r$ is the reflected Stokes vector.

$$\boldsymbol{S}_r = \boldsymbol{M}\boldsymbol{S}_i \tag{2.26}$$

This matrix can be used to find a reflected polarization state for a known situation given a set of parameters for the material. The main polarimetric component of the Mueller matrix comes form the Fresnel reflectance coefficients presented in the next section.

2.5.2 Fresnel Reflectance Equations . The Fresnel reflectance equations define the amount of in-plane and out-of-plane reflection, r_p and r_s, respectively. The parameters of importance for these equations include the complex index of refractions, $\tilde{n}$, of both the object and the air, and the pitch of the surface relative to the incident light, θ_i. The following equations describe these relationships.

$$r_s(\theta_i) = \frac{2\tilde{n}_i \cos\theta_i}{\tilde{n}_i \cos\theta_i + \sqrt{\tilde{n}_r^2 - \tilde{n}_i^2 \sin^2\theta_i}} \tag{2.27}$$
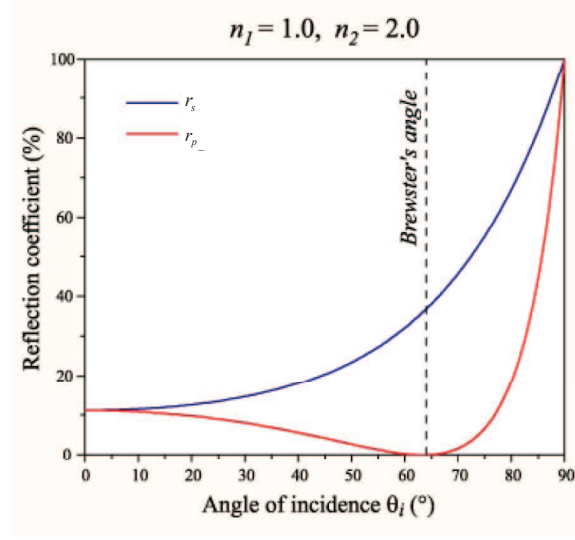
Figure 2.6:   In-plane and out-of-plane reflectance curves for a given index of refraction are computed from Fresnel equations. The difference between these curves constitutes the degree of polarization [5].

$$r_p(\theta_i) = \frac{2\tilde{n}_i \tilde{n}_r \cos\theta_i}{\tilde{n}_r^2 \cos\theta_i + \tilde{n}_i \sqrt{\tilde{n}_r^2 - \tilde{n}_i^2 \sin^2\theta_i}} \tag{2.28}$$

In these equations, $\tilde{n}_i$ represents the index of refraction of the medium containing the light incident on the surface, in most cases air, $\tilde{n}_r$ is the index of refraction of the surface of reflection and θ_i is the angle of light incident on the surface [14]. Figure 2.6 shows a graphical representation of these equations for a given index of refraction pair.

2.5.3   Bidirectional Reflectance Distribution Function .   For a given irradiance orientation, the Bidirectional Reflectance Distribution Function describes what fraction of the incident irradiance will be reflected into a solid angle within the hemisphere above the surface [14]. It is defined as the ratio of scattered radiance to the incident irradiance, and is a function of two dimensional angles, θ_i, θ_r and $\Delta\phi$, and the intrinsic surface parameters. The incidence angle, θ_i, is the angle between the ray of incident light and the surface normal. The angle between the reflected ray and the
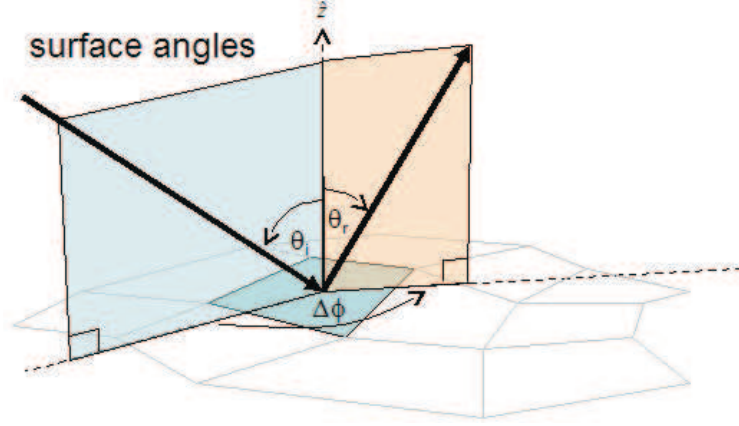
Figure 2.7: Micro-facet Model Large Scale Geometry Example. The 2-Dimensional Angles θ_i, θ_r and $\Delta\phi$ relate the incident and reflected light off of the macro-surface normal [31].

surface normal is θ_r. The angle between those two rays projected along the surface is $\Delta\phi$. Figure 2.7 shows an illustration of these angles.

Figure 2.8 shows typical BRDF models for white and black paint samples. Notice how the small highly specular spike in the white paint surface correlates to the polarization signature of the surface. The same is true but less obvious for the black paint sample.

2.5.4 Micro-facet Model . Many BRDF models may be segregated into components which represent specular scattering and volumetric scattering. Figure 2.9 shows a representation of these types of scattering.

Priest and Germer further decomposed these models into a micro-facet representation [28]. The micro-facet representation treats the specular scattering as the result of the orientation of individual small facets on a material surface. Figure 2.10 shows an illustration of the relation of the micro-surface to the macro-surface. The decomposition of a BRDF model into the micro-facet representation thus enables polarization of the model via the Fresnel reflectance off each individual micro-facet. In
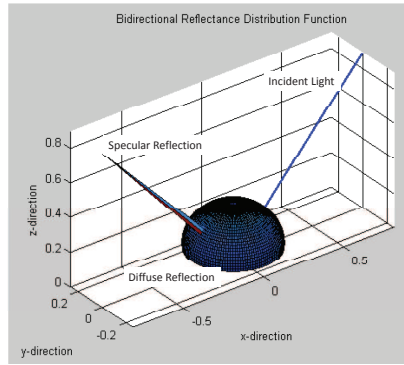
(a) White Paint BRDF
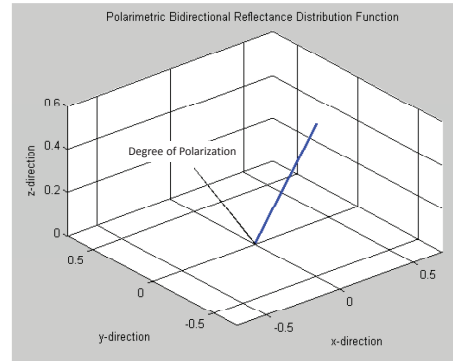
(b) White Paint directional DoLP
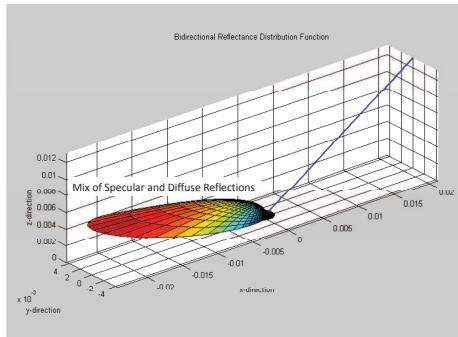
(c) Black Paint BRDF
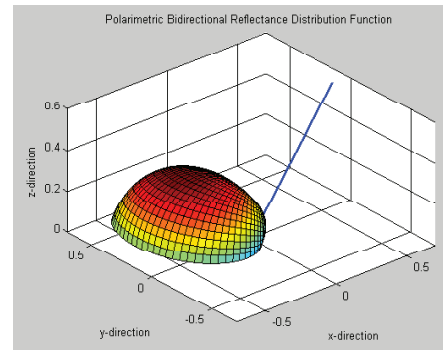
(d) Black Paint directional DoLP

Figure 2.8:   Bidirectional Reflectance Distribution Function Examples. Notice how the small, highly specular spike in the white paint surface correlates to the polarization signature of the surface. The same is true but less obvious for the black paint sample.
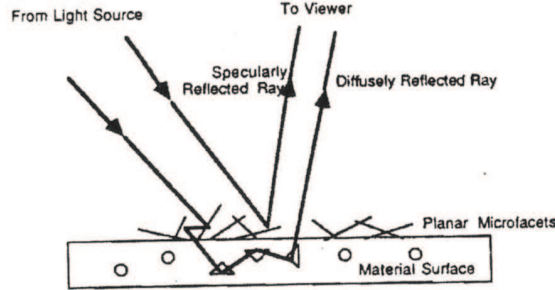
Figure 2.9: Example of Types of Scattering. There are two major types of scattering. Specular from planar surfaces and diffuse or volumetric scattering from multiply reflected components. Specular reflections contribute to the polarization state of reflected light while diffuse components have no polarization [29].

addition, the volumetric scattering component is usually considered to be additive and completely unpolarized.

Each micro-facet is considered to be oriented at an angle θ_N relative to the macro-surface normal. Half the angle between the source and the receiver is known as the bistatic angle, β. Equation (2.29) shows how to obtain β given θ_i, θ_r and ϕ. Figure 2.10 illustrates these angles.

$$\beta = \frac{1}{2} \cos^{-1}[\cos \theta_i \cos \theta_r + \sin \theta_i \sin \theta_r \cos \phi] \tag{2.29}$$

Given β, θ_N can be determined as

$$\theta_N = \cos^{-1}\left[\frac{\cos \theta_i + \cos \theta_r}{2 \cos \beta}\right] \tag{2.30}$$

2.5.4.1 Jones Matrix . A Jones matrix is an adequate means of transferring polarized energy when only Fresnel reflection is considered. The Jones matrix transforms the incident electric field oriented in the s and p polarization states to the reflected s and p polarization states. This matrix transformation is given by
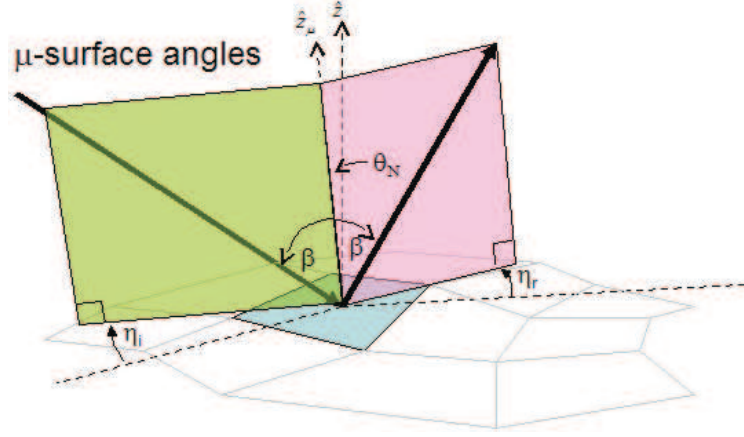
Figure 2.10: Geometry of the Micro-Surface Relative to the Macro-Surface. The offset of a given micro-surface is defined by the angle between the normals, θ_N, and the rotation angles for incident and reflected light, η_i and η_r. The angle β is the half angle between the incident and reflected light [31].

$$
\begin{bmatrix} E_s^r \\ E_p^r \end{bmatrix} = \begin{bmatrix} r_s & 0 \\ 0 & r_p \end{bmatrix} \begin{bmatrix} E_s^i \\ E_p^i \end{bmatrix}
\tag{2.31}
$$

where E_s^i and E_p^i are the magnitudes of the incident s and p polarization electric field, r_s and r_p are the Fresnel coefficients, and E_s^r and E_p^r are the reflected s and p polarization electric field magnitudes. This is equivalent to the Fresnel reflectance equation, if the incident light is considered to be unpolarized, or the magnitudes of the incident electric fields are equal.

2.5.4.2 Jones Matrix to Mueller Matrix Conversion. Given a Jones matrix, an equivalent Mueller matrix may be developed, although the converse is not true since a Mueller matrix handles the more general case of volumetric scattering.

It can be seen from Figure 2.10, two coordinate transformations are required to maintain Poynting vectors in the same frame of reference. The first transformation

rotates the incident light vector from the plane of reflectance to the plane of incidence about the macro-surface normal and is given by η_i. The second transformation rotates the specular plane of incidence to the plane of reflectance about the micro-facet surface normal and is given by η_r.

$$\cos(\eta_i) = \frac{\frac{\cos(\theta_i)+\cos(\theta_r)}{2\cos(\beta)} - \cos(\theta_i)\cos(\beta)}{\sin(\theta_i)\sin(\beta)} \tag{2.32}$$

$$\cos(\eta_r) = \frac{\frac{\cos(\theta_i)+\cos(\theta_r)}{2\cos(\beta)} - \cos(\theta_r)\cos(\beta)}{\sin(\theta_r)\sin(\beta)} \tag{2.33}$$

The coordinate transformation of the electric fields is accomplished by multiplying the incident electric field, which is defined in terms of the macro-facet coordinate system relative to the incident illumination direction, by the η_i coordinate transformation before the Fresnel reflectance. After the Fresnel reflectance, the η_r coordinate transformation is accomplished, which produces the reflected electric field components in terms of the sensor coordinate system. The resulting Jones matrix is given by

$$\begin{bmatrix} E_s^r \\ E_p^r \end{bmatrix} = \begin{bmatrix} \cos(\eta_r) & \sin(\eta_r) \\ -\sin(\eta_r) & \cos(\eta_r) \end{bmatrix} \begin{bmatrix} r_s & 0 \\ 0 & r_p \end{bmatrix} \begin{bmatrix} \cos(\eta_i) & -\sin(\eta_i) \\ \sin(\eta_i) & \cos(\eta_i) \end{bmatrix} \begin{bmatrix} E_s^i \\ E_p^i \end{bmatrix} \tag{2.34}$$

$$\begin{bmatrix} E_s^r \\ E_p^r \end{bmatrix} = \begin{bmatrix} T_{ss} & T_{ps} \\ T_{sp} & T_{pp} \end{bmatrix} \begin{bmatrix} E_s^i \\ E_p^i \end{bmatrix} \tag{2.35}$$

Now the Jones matrix components are used to construct the Fresnel reflectance Mueller matrix. The complete 4×4 Mueller matrix may be reproduced from these elements, but only the 3×3 matrix components relevant to linear polarization states are shown. The elements of the specular component of the Mueller matrix, R_{Fxx}, expressed in terms of the Jones matrix components and their complex conjugates, represented by a superscript asterisk, are given as [28]:

27

$$R_{F11} = \frac{1}{2}[|T_{ss}|^2 + |T_{sp}|^2 + |T_{ps}|^2 + |T_{pp}|^2] \tag{2.36}$$

$$R_{F12} = \frac{1}{2}[|T_{ss}|^2 + |T_{sp}|^2 - |T_{ps}|^2 - |T_{pp}|^2] \tag{2.37}$$

$$R_{F13} = \frac{1}{2}[T_{ss}T_{ps}^* + T_{ss}^*T_{ps} + T_{sp}T_{pp}^* + T_{sp}T_{pp}^*] \tag{2.38}$$

$$R_{F21} = \frac{1}{2}[|T_{ss}|^2 - |T_{sp}|^2 + |T_{ps}|^2 - |T_{pp}|^2] \tag{2.39}$$

$$R_{F22} = \frac{1}{2}[|T_{ss}|^2 - |T_{sp}|^2 - |T_{ps}|^2 + |T_{pp}|^2] \tag{2.40}$$

$$R_{F23} = \frac{1}{2}[(T_{ss}T_{ps}^* + T_{ss}^*T_{ps}) - (T_{ps}T_{pp}^* + T_{ps}T_{pp}^*)] \tag{2.41}$$

$$R_{F31} = \frac{1}{2}[T_{ss}T_{sp}^* + T_{ss}^*T_{sp} + T_{ps}T_{pp}^* + T_{ps}T_{pp}^*] \tag{2.42}$$

$$R_{F32} = \frac{1}{2}[(T_{ss}T_{sp}^* + T_{ss}^*T_{sp}) - (T_{ps}T_{pp}^* + T_{ps}T_{pp}^*)] \tag{2.43}$$

$$R_{F33} = \frac{1}{2}[(T_{ss}T_{pp}^* + T_{ss}^*T_{pp}) - (T_{ps}T_{sp}^* + T_{ps}T_{sp}^*)] \tag{2.44}$$

$$\tag{2.45}$$

Torrance and Sparrow presented one of the first polarimetric BRDF, pBRDF, models to capture the off-specular peak and to provide better predictions as θ_r becomes more glancing [33]. Maxwell and Beard improved upon this model in order to better represent paint samples [23], and Shell furthered the model to incorporate more target like materials [31]. The next three sections describe the differences between these models and why the Shell model was chosen for this work.

2.5.4.3 Torrance-Sparrow Model. Torrance and Sparrow treat each micro-facet as a specular surface for which the surface normal angular positions, α, are distributed along a Gaussian function, $P(\alpha)$. The diffuse component of the BRDF arises from multiple micro-facet reflections or internal scattering as seen in Figure 2.9. The reflected radiance, L_r, is then expressed as the sum of the specular and diffuse components.

$$L_r = L_{r,s} + L_{r,d} \tag{2.46}$$

The diffuse component is given in terms of the incident radiance by the Lambertian reflectance equation

$$L_{r,d} = aL_i \cos(\theta_i) \tag{2.47}$$

where a is a fitting constant, L_i is the radiance of the incident light and θ_i is the angle of incident light.

The specular reflection is obtained by estimating the Fresnel reflection from each micro-facet. The significant advancement made from this model was the introduction of an attenuation factor, G, which accounts for masking and shadowing in the micro-facet surface. Masking is the blockage of specular reflections by adjacent micro-facets, while shadowing is the blockage of the illumination source to one micro-facet by adjacent micro-facets. The resulting BRDF from the Torrance-Sparrow model is

$$\boldsymbol{M}_r = \boldsymbol{R}_F(\theta_i, \hat{n})P(\alpha)G + \boldsymbol{L}_{r,d} \tag{2.48}$$

where $\boldsymbol{R}_F(\theta_i, \hat{n})$ is the Fresnel reflectance associated with incident light. The distribution of the micro-facets is described by the probability density function given in Equation (2.49). A roughness parameter c that relates the distribution of the facet slopes relative to the normal plane is required. The parameter α is a free fit parameter for the function and is fit empirically to the surface.

$$P(\alpha) = ce^{-c^2\alpha^2} \tag{2.49}$$

Torrance and Sparrow use a value of $c = 0.05$, which was justified based on fitting the data to experimentally determined BRDFs [33]. The shadowing parameter used by Torrance and Sparrow, G, is given in Equation (2.50). The G component

29

describes an attenuation factor for the specular component. Therefore, when $G = 1$ there is no attenuation. The m and l parameters are functions of the incidence angle and are free fit for a particular surface.

$$G(\theta_i) = 1 - \frac{m(\theta_i)}{l(\theta_i)} \tag{2.50}$$

While the Torrance-Sparrow model makes use of first principles to model the BRDF, it nonetheless requires parameters to be fit to experimental data. Modifications to the Torrence-Sparrow model by Maxwell and Beard are covered in the next section.

2.5.4.4 Maxwell-Beard Model. The Maxwell-Beard BRDF model was originally developed for use on painted surfaces [23]. As with the Torrance-Sparrow model, specular and diffuse contributions to the BRDF are considered separately. The complete Maxwell-Beard BRDF model is given by the sum of the surface and volumetric components given in Equations (2.52) and (2.54), or

$$\boldsymbol{M}_r(\theta_i, \phi_i; \theta_r, \phi_r) = \boldsymbol{M}_{r_{spec}} + \boldsymbol{M}_{r_{dif}} \tag{2.51}$$

With the Maxwell-Beard model, only single reflections from the micro-facet surface are considered. The specular component of the BRDF may be expressed as

$$\boldsymbol{M}_{r_{spec}}(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\boldsymbol{R}_F(\beta)}{\boldsymbol{R}_F(0)} \frac{f_{ZBS}(\theta_N) \cos^2 \theta_N}{\cos \theta_i \cos \theta_r} SO(\tau, \Omega) \tag{2.52}$$

where β is the half angle between the source and the receiver, $\boldsymbol{R}_F$ is the Fresnel reflectance equation, expressed in terms of β, and $SO(\tau, \Omega)$ represents the shadowing function, presented in equation (2.53). The distribution of the micro-facets is obtained through a zero angle bistatic scan, f_{ZBS}, in which the detector and illumination source are co-located, or as close to the same position as possible without subtending each other. The surface normals of each micro-facet are defined as being oriented in the

(θ_N, ϕ_N) direction. The measured signal of the ZBS scan is related to the density of micro-facets. As in the Torrance-Sparrow model, reflections from the micro-facets are given by the Fresnel reflectance equations.

All of the parameters needed for the Maxwell-Beard BRDF have to be assumed or obtained experimentally. The Fresnel reflectance requires the complex index of refraction $\tilde{n}$ of the material. Maxwell and Beard were able to make some assumptions in their studies. They assumed that the surfaces were dielectrics, which allowed them to consider $k = 0$ or $\tilde{n} = n$. A value of n in their study was estimated as $n = 1.65$ and was based on experience with paint samples [23]. As an alternative, Maxwell and Beard indicate the value of n may be calculated based on Brewster's angle.

Maxwell and Beard found similar variations caused by shadowing and masking of the micro-facets, previously addressed in the Torrance-Sparrow model discussion. However, they developed their own empirically derived function to account for shadowing and obscuration, SO, which they found superior to the Torrance-Sparrow function. The SO function has two free parameters, τ and Ω, and is given by

$$SO(\tau, \Omega) = \frac{1 + \frac{\theta_N}{\Omega} e^{-2\beta/\tau}}{1 + \frac{\theta_N}{\Omega}} \left(\frac{1}{1 + \frac{\phi_N \theta_i}{\Omega^2}} \right) \tag{2.53}$$

where ϕ_N adjusts the falloff rate of the shadowing and obscuration function.

The non-Lambertian volume component development was motivated by experimental observation that the diffuse scatter component was in fact not Lambertian due to both the angular dependency and the lack of complete depolarization. Derivation of this volume component considers the exponential loss via scattering of energy as the light propagates into the medium as well as the exponential loss of energy as the light propagates back to the surface. It is assumed that there is no net transmission of energy through the surface, and absorption in the medium is not explicitly considered. Given these considerations, the diffuse component of the BRDF is given as

$$\boldsymbol{M}_{r_{dif}} = \frac{2\rho_v f(\beta) g(\theta_N)}{\cos\theta_i + \cos\theta_r} \tag{2.54}$$

where $f(\beta)$ and $g(\theta_N)$ include the β and θ_N dependencies and are treated as free parameters for adjustment based on the empirical data. However, the model implemented by Maxwell and Beard kept $f(\beta) = h(\theta_N) = 1$ and simply states that these parameters may provide flexibility in future model development [23]. ρ_v is experimentally obtained by measuring the BRDF at $\theta_i = \theta_r = 0°$ with the incident light polarized orthogonally to the linear polarizing filter.

2.5.5 *Shell Target Model* . The basis of the Shell Target model follows from the polarization of micro-facet BRDF models presented above. It is similar to each of these models in that it has been decomposed into contributions from the specular and volumetric components. The differences in this model are due to the makeup of each of these components and the assumptions used to fit a model to a more broad 'target' criterion which makes it ideal for the wide verity of man-made targets found in an indoor environment. Another attraction to this model is that an extensive database of materials exists with these model parameters. The National Geospatial Intelligence Agency's Nonconventional Exploitation Factors Database (NEFDS) contains parameters for the Maxwell-Beard model which may be polarized by the application of the Priest-Germer micro-facet polarization technique.

The Shell Target model can be broken into four sub-equations consisting of the Fresnel reflectance off a micro-facet, $\boldsymbol{R}_F(\beta)$, given in Equation (2.45), the probability density function of the orientation of micro-facets, $p(\theta_N)$, the shadowing term, SO, and the diffuse component, $\boldsymbol{M}_{dif}$.

The effect of $p(\theta_N)$ is to place a 'peak' in the specular direction. This model can use one of two different fit functions for the micro-facet probability distribution function, $p(\theta_N)$, a Gaussian fit as in the Torrance-Sparrow model and a Modified Cauchy fit. The micro-facet probability distribution function may be thought of as providing

32

the 'spread' of the specular micro-facet reflections according to the surface roughness statistics. This distribution function uses two parameters: a surface roughness parameter, σ, and a bias parameter, B. A smaller σ corresponds to a smoother or more specular surface. The B parameter provides an overall magnitude adjustment.

The Gaussian micro-facet probability distribution function $p_G(\theta_N)$ is used by the Priest-Germer model and is given by Equation (2.55). The Modified Cauchy probability distribution function $p_C(\theta_N)$ is adapted from that used by the more recent versions of the NEF Maxwell-Beard BRDF model and is shown in Equation (2.56).

$$p_G(\theta_N) = \frac{B}{2\pi\sigma^2 \cos^2(\theta_N)} e^{\frac{-\tan^2(\theta_N)}{2\sigma^2}} \tag{2.55}$$

$$p_C(\theta_N) = \frac{B}{\cos(\theta_N)(\sigma^2 + \tan^2(\theta_N))} \tag{2.56}$$

The shadowing function is meant to factor in shadows caused by surface roughness and glancing angles. The Shell Target model uses a simplified version of the Maxwell-Beard Shadowing function, given by

$$SO(\theta_N) = \frac{1 + \frac{\theta_N}{\Omega} e^{-1\beta/\tau}}{1 + \frac{\theta_N}{\Omega}} \tag{2.57}$$

Finally, the diffuse scattering component of the generalized micro-facet model is considered. The only representation given is that from the NEF Maxwell-Beard BRDF model. This term is completely randomly polarized and is expressed as

$$\boldsymbol{M}_{dif} = \rho_D + \frac{2\rho_v}{\cos(\theta_i) + \cos(\theta_r)} \tag{2.58}$$

where ρ_D and ρ_V are two empirical fit parameters. The nomenclature used is identical to the NEF, and is an adaptation of that originally proposed by Maxwell, where ρ_D is the diffuse or Lambertian component and ρ_V is the volumetric scattering parameter.

The Maxwell-Beard model does allow for non-physical, negative values of ρ_D and ρ_V, in order to provide a better fit to the empirical data.

As with each of the other micro-facet models, the resulting Mueller matrix may be expressed as the sum of the specular and diffuse components.

$$\boldsymbol{M} = \boldsymbol{R}_F(\beta)p(\theta_N)S + \boldsymbol{F}_{vol} \qquad (2.59)$$

In general the surface parameters for the Shell Target model are fit empirically and are unknown for a particular surface. However, some simplifications can be made in the equations in order to better estimate surface geometry. Chapter IV, Section 4.3 describes a set of simplifying assumptions that are common in indoor environments and may be used to better estimate the geometry of a scene.

The broad range of materials that work with the Shell model along with the availability of an extensive list of surfaces and the fact that this particular model is used in the simulation software presented in Chapter III make this the ideal model to use for this research.

2.6 Polarimetric Shape Recovery

The first paper in computer vision to use polarization information was by Koshikawa, who used an ellipsometry technique to constrain surface normals on dielectric surfaces [18]. Koshikawa used a polarization reflectance model based on the Mueller calculus for specular reflection and the Stokes vector representation for polarization. Two other sets of researches followed in similar work but used different models or constraints for their work. The first was Lawrence Wolff who's work is described in Section 2.6.1. A significantly different approach was taken by Pablo d'Angelo and Christian Wohler and is described in Section 2.6.2.

2.6.1 Lawrence Wolff .   Wolff developed a polarization reflectance model that is called the Fresnel reflectance model because it is a geometric reflectance model

that utilized the Fresnel reflection coefficients directly [34,35]. The Fresnel reflectance method is similar to the degree and angle of polarization equations, but bypasses the Stokes vector math by fitting a sine function of frequency two hertz to the intensity measurements as a function of polarization angle. Equation (2.60) shows the equation of the Fresnel reflection ratio.

$$F_r = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} \qquad (2.60)$$

The angle of polarization is then simply the phase offset of the sine wave. Wolff uses this ratio to segregate materials in an image into dielectric and metal components [36,38].

Another useful technique developed by Wolff to determine surface structure was a binocular polarization-based technique which determines surface orientation from the intersection of two specular planes, each constraining the surface normal in two of three dimensions [37].

The advantage of the binocular method over a monocular polarization-based methods for unique determination of surface orientation is that no knowledge of the index of refraction for the material surface is required. The disadvantage is that points need to be corresponded between images. This is less of a problem when determining surface orientation for a flat surface where the intersection of any two specular planes determined from two points with the same surface orientation will compute the correct surface orientation. However, smooth curved surfaces without any distinctive markings pose a problem.

Wolff suggests that his binocular method could be very useful in conjunction with other surface correspondence techniques such as depth from shading or depth from defocus [37].

2.6.2   d'Angelo and Wohler .   Pablo d'Angelo and Christian Wohler use an analysis of reflectance and polarization properties to reconstruct rough metallic

surfaces [10–12]. These surfaces are regarded as being composed of micro-facets of random orientation. However, d'Angelo and Wohler implement their own empirically fit functions for intensity and degree of polarization.

They create models for intensity, degree of linear polarization and angle of polarization measurements to fit to the rough metallic surfaces with which they are working. Their method of photo-stereo imaging relies on a pair of polarization images in which d'Angelo and Wohler make the assumption that the scene is illuminated by unpolarized point light sources at known locations. They then use a Levenburgh-Marquardt algorithm to estimate surface gradients.

Each of these methods has advantages and disadvantages. Some require a structured environment or stationary image frames while others require model parameters to be fit for a specific material. In order to produce the most realistic results, this paper proposes the use of the Shell Target model as the primary non-linear model for simulation and estimation.

The Shell Target model will be used to estimate camera attitude. These attitude states are maintained and updated through a Kalman filter approach. The next section described the Unscented Kalaman filter used in this research.

2.7   The Unscented Kalman Filter

A Kalman filtering technique is used to maintain estimates of the state of a system as sets of measurements become available [24]. For the problem of navigation, uncertainty in position, velocity, and attitude grow unless constrained by stable measurements. The measurement used in this thesis constrains the attitude estimation by providing periodic measurements of relative angles between the camera and interior surfaces. This section explains a basic background needed to understand the uncertainty in state estimation and how the measurement update process improves the uncertainty.

One algorithm for solving the problem of non-Gaussian, non-linear filtering is the extended Kalman filter [3]. This filter is based on the idea of linearizing the required models. However, these approximations can lead to poor representations of non-linear functions and cause divergence. Another algorithm for generalized filtering, known as a particle filter, is performed by using a set of Monte Carlo simulations [25]. This type of method allows for a complete representation of the posterior distribution of the states, so that a full statistical distribution can be computed. However, this method requires a great computational capability.

Julier and Uhlmann present the unscented Kalman filter as a method of applying Kalman filtering techniques to non-linear systems [3]. The unscented Kalman filter addresses some of the approximation issues of the extended Kalman filter without the computational requirements of the particle filter. The UKF uses the set of true non-linear models for state propagation and measurement updates and approximates the distribution of the states by a Gaussian random variable. This section explains the basics of the unscented Kalman filter by first describing the unscented transform, in Section 2.7.1, and then the presenting the Kalman filter update process, in Section 2.7.2.

2.7.1 Unscented Transform .   The unscented transform is a way of calculating the statistics of a random variable after a non-linear transformation. In the case of this thesis, the non-linear transformation converts the state estimate to a measurement prediction through the set of measurement equations presented in Chapter IV. A set of state variables, $\boldsymbol{x}$, with mean $\bar{\boldsymbol{x}}$ and covariance $\boldsymbol{P}_{xx}$, is passed through the non-linear equation, $\boldsymbol{y} = h(\boldsymbol{x})$. The random variable, $\boldsymbol{y}$, then has mean $\bar{\boldsymbol{y}}$ and covariance $\boldsymbol{P}_{yy}$, which can be found using the unscented transform.

Given a state vector of length n, the first step of this process is to create a $2n + 1$ set of weighted sigma points such that they capture the mean and covariance of the prior random variable. In order to better capture higher order terms of the probability density function, and to maintain a positive definite state uncertainty, the

scaling parameters κ and α must be chosen properly. The κ factor is set to be greater than or equal to 0, so that the state uncertainty will remain positive definite. The α weighting factor determines the extent of spread for the sigma points. Reference [25] for more information on the best choices for these terms for a particular problem. The scaling parameter λ can be found using Equation (2.61).

$$\lambda = \alpha^2(n + \kappa) - n \tag{2.61}$$

The sigma point selection and scaling can then be performed through

$$\boldsymbol{X}_0 = \bar{\boldsymbol{x}} \tag{2.62}$$

$$\boldsymbol{X}_i = \bar{\boldsymbol{x}} + (\sqrt{(n+\lambda)\boldsymbol{P}_{xx}})_i \qquad i = 1, \ldots, n \tag{2.63}$$

$$\boldsymbol{X}_i = \bar{\boldsymbol{x}} - (\sqrt{(n+\lambda)\boldsymbol{P}_{xx}})_i \qquad i = n+1, \ldots, 2n \tag{2.64}$$

$$W_0^{(m)} = \lambda/(n+\lambda) \tag{2.65}$$

$$W_0^{(c)} = \lambda/(n+\lambda) + (1 - \alpha^2 - \beta) \tag{2.66}$$

$$W_i^{(m)} = W_i^{(c)} = 1/2(n+\lambda) \qquad i = 1, \ldots, 2n \tag{2.67}$$

where $\boldsymbol{X}_0$ is the mean sigma point, $\boldsymbol{X}_i$ are the additional sigma points and $W_i^{(m)}$ and $W_i^{(c)}$ are the weights associated with those points, used to determine the mean and covariance, respectively. The weighting of the zeroth sigma point affects errors in higher order terms of the probability density function. The β term is a weighting term which allows for minimization of these errors if prior knowledge of the distribution is available. For the Gaussian distributions assumed throughout this thesis, a value of $\beta = 2$ is used.

These sigma points are propagated through the true non-linear transformation, such that $\boldsymbol{Y}_i = h(\boldsymbol{X}_i)$. Then, the mean and covariance of the transformed sigma points are calculated using Equations (2.68) - (2.69).

$$\bar{\boldsymbol{y}} = \sum_{i=0}^{2n_x} W_i^{(m)} \boldsymbol{Y}_i \tag{2.68}$$

$$\boldsymbol{P}_{yy} = \sum_{i=0}^{2n_x} W_i^{(c)} (\boldsymbol{Y}_i - \bar{\boldsymbol{y}})(\boldsymbol{Y}_i - \bar{\boldsymbol{y}})^T \tag{2.69}$$

The information presented in this subsection is a general form of the unscented transform. The specific form associated with the unscented Kalman filter is presented next.

2.7.2 Kalman Filter Process . The general Kalman filter process can be broken down into a state transition and state measurement model. In this case, the set of nonlinear equations has the general form of

$$\boldsymbol{x}(t_i^-) = f(\boldsymbol{x}(t_{i-1}^+), \boldsymbol{v}(t_{i-1})) \tag{2.70}$$

$$\boldsymbol{z}(t_i) = h(\boldsymbol{x}(t_i^-), \boldsymbol{n}(t_i)) \tag{2.71}$$

where $\boldsymbol{x}(t_i^-)$ is the current state of the system just before a measurement update, $\boldsymbol{v}(t_i)$ is the propagation noise, $\boldsymbol{z}(t_i)$ is the set of measurement observations and $\boldsymbol{n}(t_i)$ is the measurement noise.

The work presented in this thesis is primarily concerned with the state measurement model. More information on the unscented Kalman filter state transition algorithm can be found in [3]. For the work presented in this thesis, an initial condition for states and uncertainty are given to represent those found just prior to a measurement.

When a measurement becomes available, the first step is to make a prediction of the measurement from the current state estimates. This is done through the unscented transform presented in the previous section. The generalization of the measurement

equations is $\boldsymbol{Z}_i(t_i) = h(\boldsymbol{X}_i(t_i^-))$. Equations (2.72) and (2.73) show the mean and uncertainty for the predicted measurement, $\bar{\boldsymbol{z}}$ and $\boldsymbol{P}_{zz}$, respectively.

$$\bar{\boldsymbol{z}}(t_i) = \sum_{i=0}^{2n} W_i^{(m)} \boldsymbol{Z}_i(t_i) \tag{2.72}$$

$$\boldsymbol{P}_{zz}(t_i) = \sum_{i=0}^{2n} W_i^{(c)} [\boldsymbol{Z}_i(t_i) - \bar{\boldsymbol{z}}(t_i)][\boldsymbol{Z}_i(t_i) - \bar{\boldsymbol{z}}(t_i)] \tag{2.73}$$

The uncertainty between the prediction of the measurement and the observed measurement, known as the residual covariance, is determined through Equation (2.74).

$$\boldsymbol{P}_{xz}(t_i) = \sum_{i=0}^{2n} W_i^{(c)} [\boldsymbol{X}_i - \bar{\boldsymbol{x}}(t_i^-)][\boldsymbol{Z}_i(t_i) - \bar{\boldsymbol{z}}(t_i)] \tag{2.74}$$

The state update is then performed by first finding the Kalman gain, $\boldsymbol{K}(t_i)$, in Equation (2.75).

$$\boldsymbol{K}(t_i) = \boldsymbol{P}_{xz}(t_i)\boldsymbol{P}_{zz}^{-1}(t_i) \tag{2.75}$$

The Kalman gain is used to weight the error between the predicted and observed measurements in order to determine the new state estimate, $\bar{\boldsymbol{x}}(t_i^+)$, in Equation (2.76). It is also used to determine the new state uncertainty, $\boldsymbol{P}_{xx}(t_i^+)$, in Equation (2.77).

$$\bar{\boldsymbol{x}}(t_i^+) = \bar{\boldsymbol{x}}(t_i^-) + \boldsymbol{K}(t_i)[\boldsymbol{z}(t_i) - \bar{\boldsymbol{z}}(t)] \tag{2.76}$$

$$\boldsymbol{P}_{xx}(t_i^+) = \boldsymbol{P}_{xx}(t_i^-) - \boldsymbol{K}(t_i)\boldsymbol{P}_{zz}(t_i)\boldsymbol{K}^T(t_i) \tag{2.77}$$

In general, at this point, the filter would continue to propagate and update as measurements become available. However, the algorithms presented in this thesis only use the measurement update component. Each time the Kalman filter is used, a state

estimate and associated uncertainty are given to the algorithm. A full Kalman filter algorithm could be developed in the future if a dynamics model could be determined, either using an inertial navigation system or a platform with predicable motion.

The information presented throughout this chapter is meant to give the reader an understanding of the basic concepts used for the research effort. Additional references have been provided to direct a reader to more in-depth information on each subject. Chapter III discusses the tools built using these basics in order to perform the tests presented in Chapter IV.

III. Builds and Simulations

Chapter II showed the sets of complex models which will be used to estimate surface geometry and camera orientation. Due to the complex nature and wide verity of parameters to estimate, it is beneficial to run simple simulations before proceeding to physical tests. Simulation software can be modified to easily test new ideas, then physical tests can be used to confirm the simulation or to show any additional anomalies.

For these reasons three sets of tools were used for this research. These tools are described in detail in this chapter. Section 3.1 shows a series of MATLAB graphical user interfaces, used for quick manipulation of algorithms and simple scenarios. The DIRSIG simulations software shown in Section 3.2 was used to form complex simulations. Finally, the physical polarimeter, used to perform real world tests, is described in Section 3.3

3.1   *MATLAB Simulation Models*

MATLAB Graphical User Interfaces, GUIs, were built because of their ease of use and quick adaptability. They are built to give a user an interactive feel of the dynamics of the Shell Target model, the observability of its parameters, and the ability of the estimation techniques. Descriptions of particularly helpful GUIs are presented in this section.

3.1.1   Single Parameter Estimation .   The GUI shown in Figure 3.1 allows a user to determine the errors between measurements using actual and estimated surface parameters. It was seen in Chapter II that the Shell Target model of the polarimetric-BRDF is a complex function of geometry and surface parameters. This GUI was developed to determine how each parameter effects the intensity, degree of polarization and angle of polarization. It also gives a user a better understanding of how estimation of a single parameter is effected by other parameters.
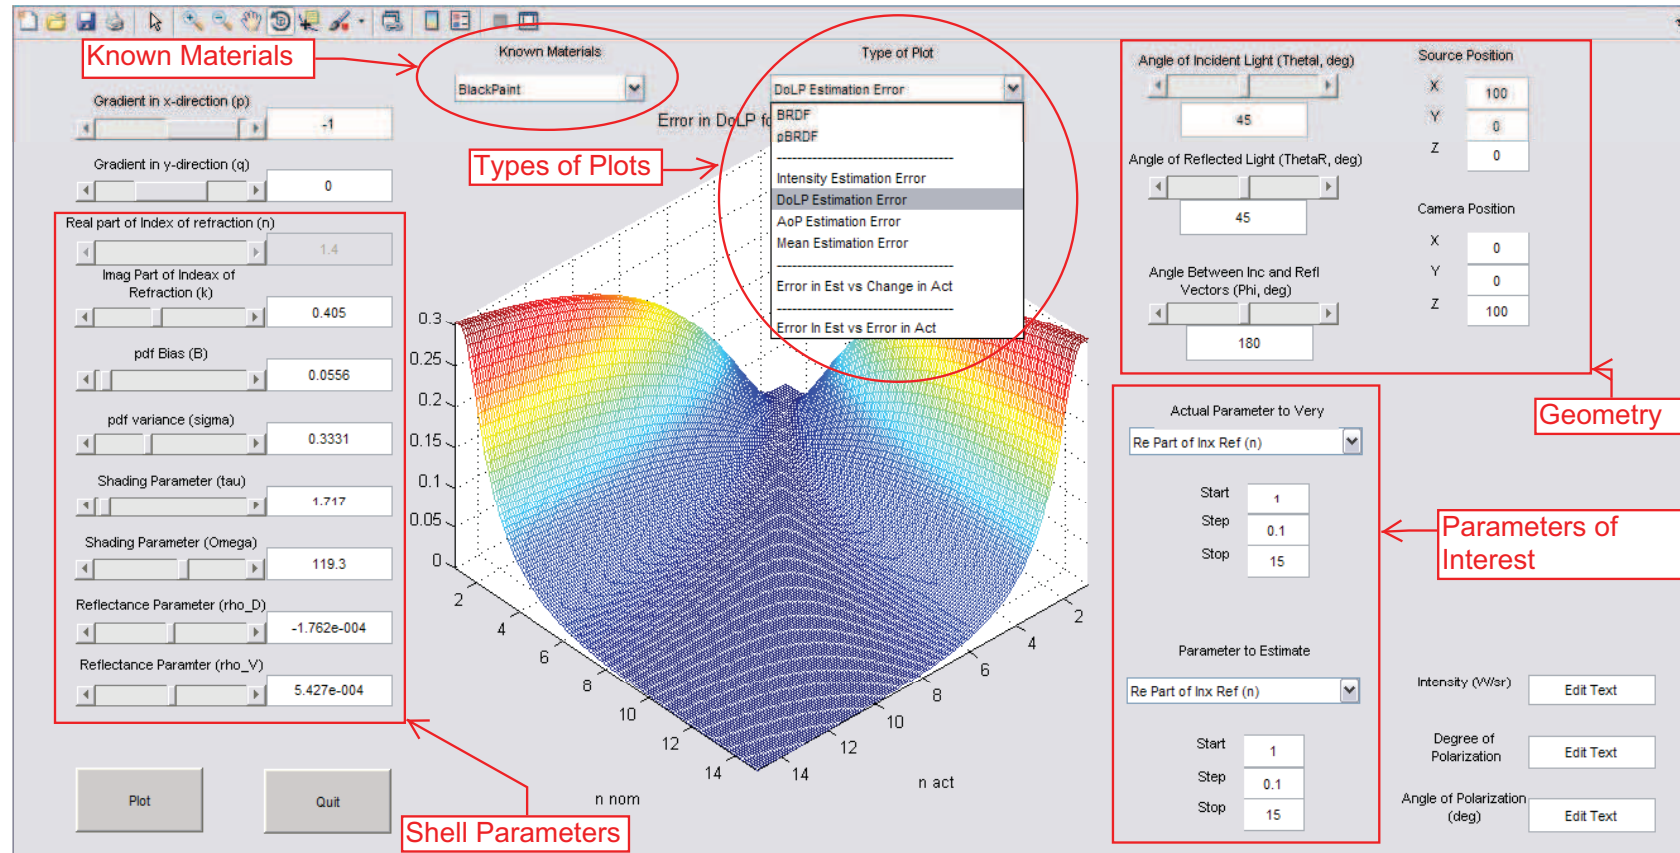
Figure 3.1:   MATLAB GUI to give a user a feel for how parameters affect the intensity, degree and angle of polarization measurements. This GUI also allows a user to determine viable constraints for estimation.

This GUI can be used to determine constraints that can be used to better estimate a set of parameters. Using an estimation technique to iteratively determine a final solution of a surface parameter given an initial estimate relies on a surface plot like the one shown in Figure 3.1. Valleys in this plot are where an estimate would eventually converge. Starting with an estimate on the wrong side of a peak in this image would lead to a wrong final estimation.

There are many settings for the user to manipulate in this GUI. The types of measurements or measurement errors are selected in the drop down menu in the top center. Once the type of plot is selected, the Shell Target parameters are chosen from the left hand side. For convenience, a drop down menu with common materials is available for use. The selection of one of these materials will automatically fill in the Shell parameters. On the right hand side of the figure, the user can define source-surface-camera geometry in two different ways, either by selecting absolute position or by selecting relative geometry. Finally, the parameters of interest for a given test are selected through the drop down menus on the lower right hand side. The user can select an actual and estimated parameter and a range of values to vary them across. Once the calculations are completed, the graph of results is displayed in the center window. This figure is a 3-dimensional representation and can be manipulated with the figure tool-bar at the top of the GUI.

3.1.2 Multiple Parameter Estimation .   Once a user has an understanding of how a single parameter can affect the measurements, the GUI shown in Figure 3.2 can be used to determine the extent to which multiple parameters may be estimated. It also allows for a user to determine observability of parameters under certain conditions and allows for determination of a viable set of constraints in order to estimate the desired parameters.

This GUI allows a user to define an actual set of surface parameters and geometry and an estimate of 'known' and unknown parameters. It allows the user to chose the parameters to estimate and even allows for multiple measurements at user

44
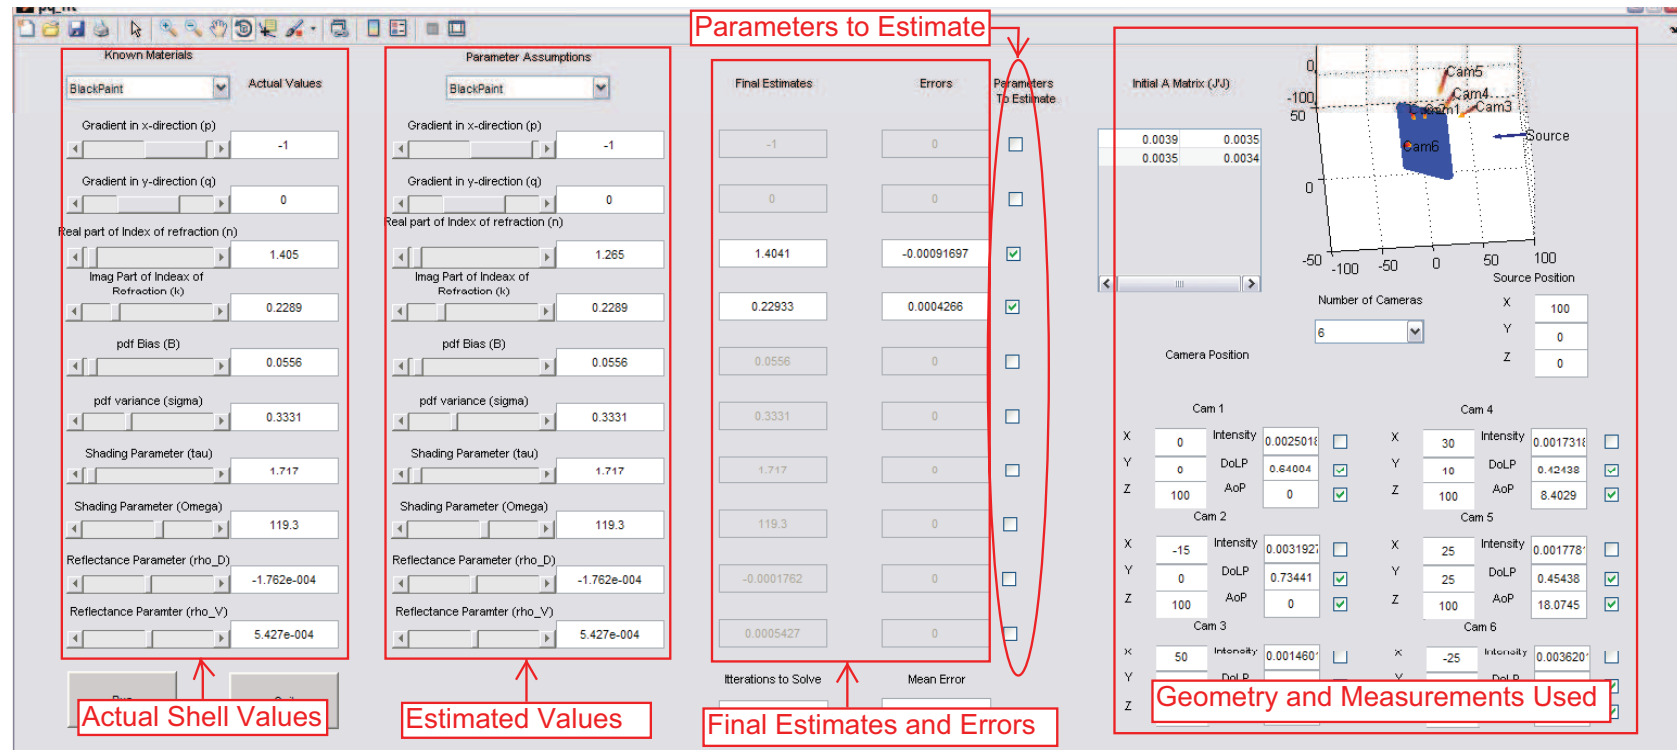
Figure 3.2: MATLAB GUI to estimate a user defined set of parameters given knowledge of the set of known parameters. This GUI allows a user to determine observability of parameters and limits of viable constraints in order to create better estimation techniques.

defined geometries. The left hand side of the GUI has the same options as the single parameter estimation GUI. It allows a user to define a full set of Shell parameters for actual and estimated values of a surface. It also includes the same drop-down menus for easily filling in the Shell parameters for known materials. A single source and up to six cameras can be placed in absolute coordinates using the options on the right hand side. The set of measurements to use can be selected with the check boxes next to each camera measurement. The set of parameters to estimate are selected by checking the boxes in the center of the figure. Once the calculations are complete, the final estimates and errors in these estimates are displayed in the center columns.

3.1.3 Multiple Hypothesis Testing .   Figure 3.3 shows a multiple hypothesis testing GUI used to return a best fit surface for a given situation from a list of known surfaces. This GUI allows a user to define actual surface parameter values and a 'known' geometry. It will then calculate errors in actual and estimated measurements and chose the material that best correlates to the measurements.

This GUI allows the user to input the actual Shell parameters in the same manner as the previous GUIs, with the values on the left hand side. A big difference in this GUI is the option to corrupt the actual surface geometry values with estimated values. This option, shown in the upper-center portion of the figure, allows a user to determine how well the geometry must be known for a given example in order to arrive at the correct solution. The same source-surface-camera geometry options are available in this GUI, noted on the right hand side of the figure. The best fit calculation is then displayed in the drop-down display in the center of the figure.

3.1.4 Estimation of beta angle GUI .   The GUI found in Figure 3.4 shows the tool used to determine how the estimation of β angle is affected by errors in estimated Shell parameters and geometry. A simplified version of this GUI is also developed using the assumptions presented in Section 4.3. The estimation of pitch, or β, angle is a main part in both the estimation of surface orientation and camera attitude estimation presented in Chapter IV.

46
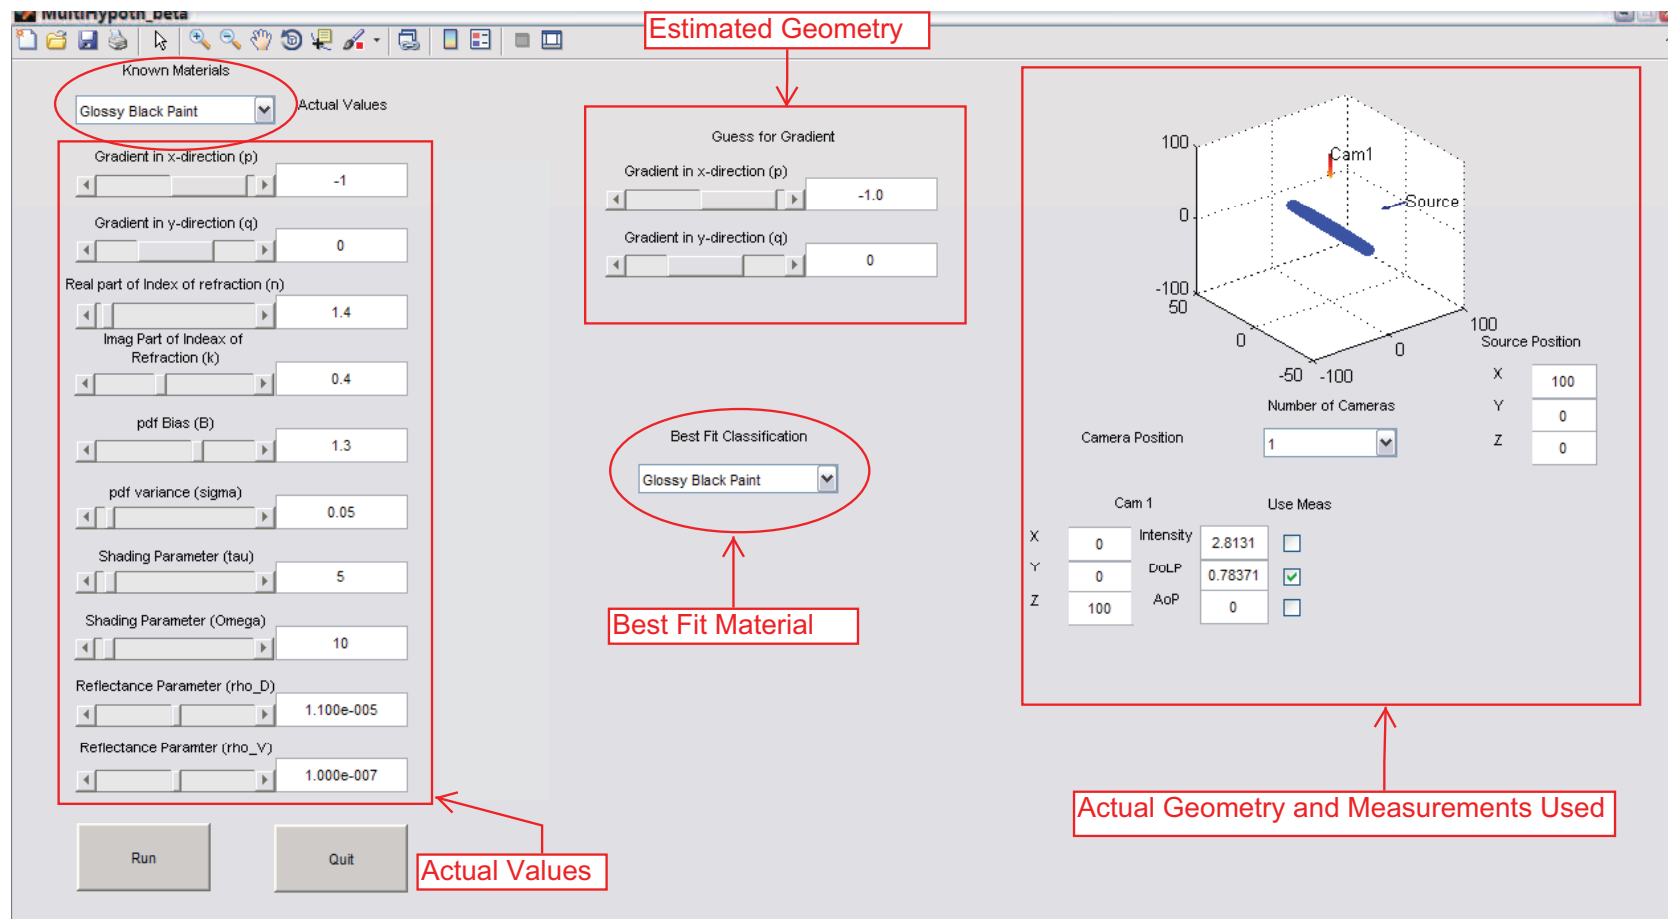
Figure 3.3: Multiple Hypothesis GUI Example. This GUI allows a user to input actual parameters and geometry for a surface. It will calculate the expected measurement for a given set of targets known to be in the scene and return the target material that most closely fits the given measurement.
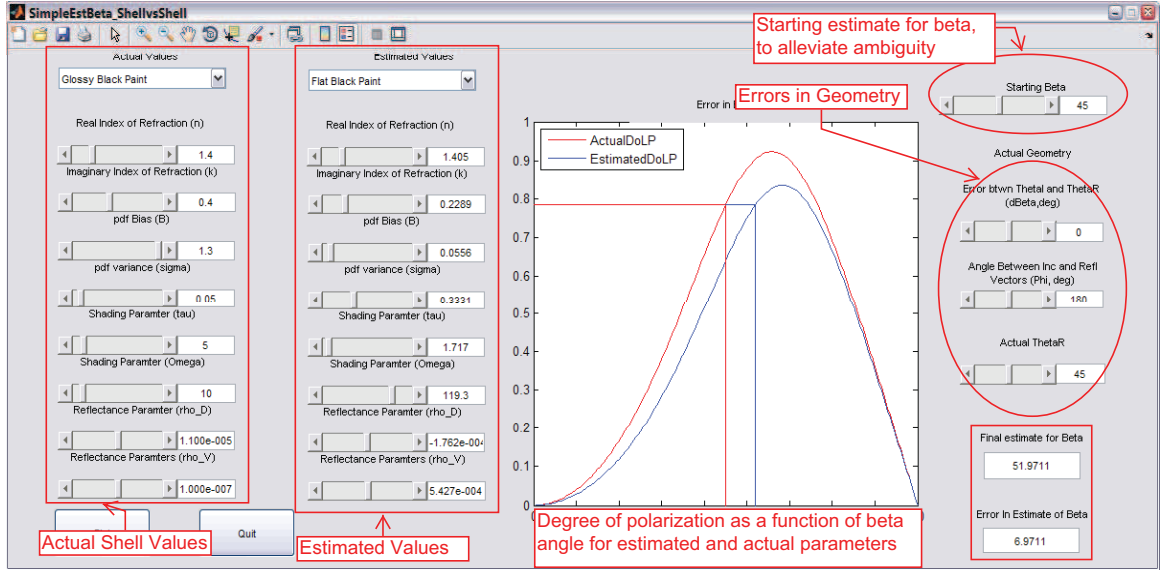
Figure 3.4:   Estimation of β angle Using Shell Parameter Values.

The left hand side of this figure allows the user to define the actual and estimated Shell parameters. In the simplified version, the estimated parameters column is replaced by estimations of the complex index of refraction only. The right hand side of the figure allows a user to choose errors in relative geometry. The starting estimation of β, chosen in the upper right hand side of the figure, accounts for the ambiguity in the estimation algorithm. Once the results are computed, a figure of the degree of polarization as a function of β angle is displayed in the center pane and the final estimation and error are given in the lower right hand portion.

3.1.5   Uncertainty in the Estimation of beta GUI .   The envelope of the uncertainty in β estimation, used in the Kalman filter approach found in Section 4.4.1, can be seen in Figure 3.5. This figure allows a user to define the actual Shell parameters for a surface and an actual geometry. Then, for descrete β angles, it will create 300 particles with random errors of up to 1% for the set of parameters. The particles' values are used to estimate a β angle, and the mean and standard deviation of the errors in these estimates are plotted in the center pane.
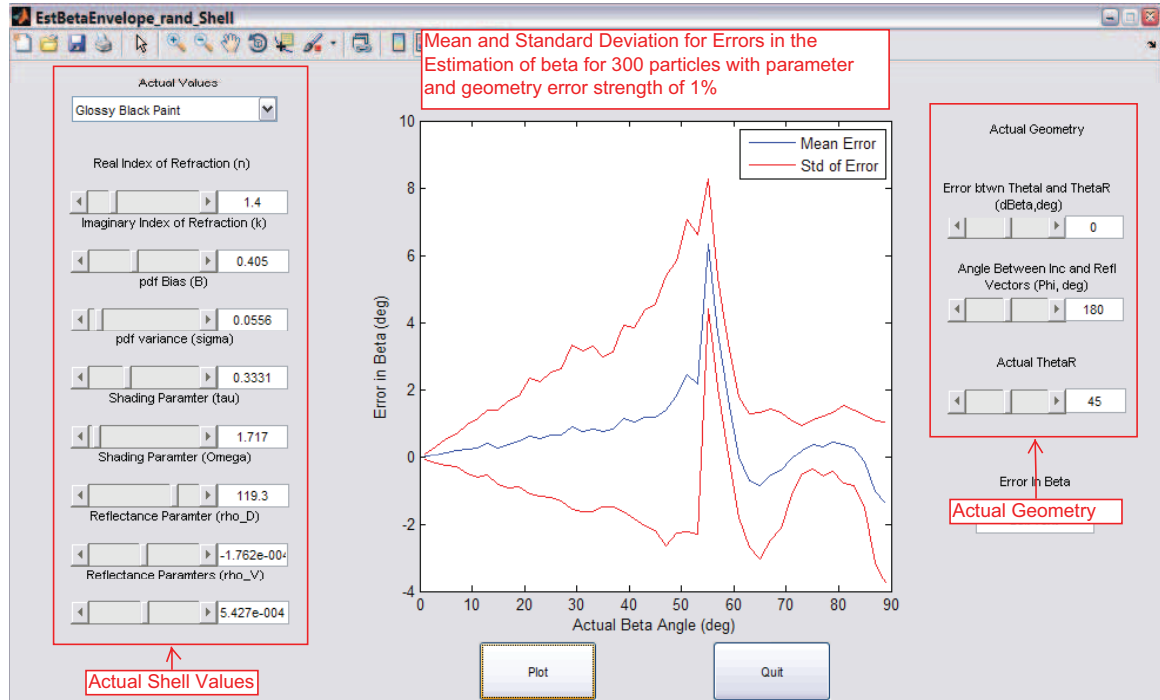
Figure 3.5:   Uncertainty in Estimation of Beta. Shell model 1% error. 300 particles. Glossy Black Paint.

This GUI allows a user to choose the actual set of Shell parameters in the same manner as the previous GUIs, by setting the individual values or choosing a material from the drop-down menu on the left hand side of the figure. The actual geometry is chosen on the right hand side of the figure.

Although MATLAB GUIs are easy to change and can be used to test new ideas quickly, they can be limited in their complexity. A simulation software package developed at the Rochester Institute of Technology has recently added a polarimetric capability. The Digital Imaging and Remote Sensing Image Generation (DIRSIG) software package was originally designed for spectral analysis of complex scenes, but recently implemented the Shell Target model into its repertoire. The next section describes the fundamentals of the DIRSIG software package.
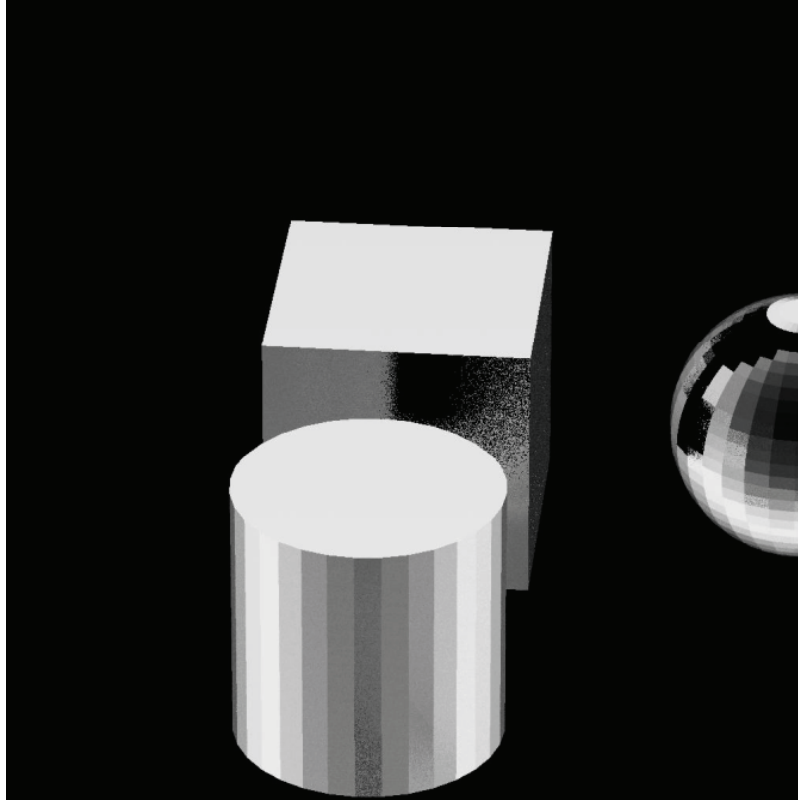
Figure 3.6: Example output image from the DIRSIG software. This image shows the intensity of three glossy black objects being illuminated by the sun.

3.2 DIRSIG Simulation Software

The DIRSIG software is a synthetic image generation model developed by the Digital Imaging and Remote Sensing Laboratory at the Rochester Institute of Technology [30]. Thanks to the work of several graduate students, this software has recently added a polarimetric imaging capability [14,26,31]. Detailed information on the inner workings of this software can be found in the literature and will not be discussed here. This section describes the simulations that were built using this software to verify the MATLAB simulations and to solidify the tests explained in Chapter IV.

A set of simple objects with known geometry was the main target for these simulations. A cube, a sphere, and a cylinder were given attributes of glossy black paint. The camera location and orientation were then set as needed to capture angles of interest. Figure 3.6 shows an example image produced from the software.

The software outputs a set of Stokes images in $W/cm^2 sr$, so that Equations (2.24) and (2.25) in Section 2.4 are all that are required to compute degree and angle of polarization images.

Atmospheric conditions were neglected in these simulations, due to the close proximity of the camera to the objects of interest, the broad spectral band of the camera and the interest in indoor environments. For images taken outdoors, this could lead to differences between the simulation and reality. Polarization of the visible spectrum from the atmosphere would be important to take into account for images taken outdoors [13]. However, for the tests in Chapter IV, small distances between the source, object and receiver mean there is little or no polarization imparted during propagation.

In order to verify the MATLAB and DIRSIG simulation models and to perform tests on more complex and realistic scenes, a physical polarimeter was built. The next section describes the construction and calibration of a visible, digital single lens refracting (dSLR) camera with a linear polarizer mounted to a rotation stage.

3.3   Physical Polarimeter

In order to validate the MATLAB and DIRSIG simulations and to determine the usefulness of the algorithms in real world scenarios, a physical polarimeter was designed and built. This important tool is instrumental in validating these simulations as well as showing additional anomalies that can arise from real world scenarios and complex conditions with multiple light sources and reflections.

3.3.1   Components .   A Sony α330 dSLR was chosen as the backbone of the physical polarimeter because of the range and controllability of its input settings. Because of the large number of camera settings and the requirement for spectral and spatial calibration for each setting, only a few complete settings were chosen for this research. The most practical setting included an F-stop of 5.7, focus at infinity, shutter speed of 1/30s, and ISO of 100. These settings were chosen because of the

typical amount of light found in an indoor environment, the desire for minimization of electrical noise between images and a long depth of focus.

The polarizer used in this system is a commercial off the shelf (COTS) item from Newport Optics. The 20LP-VIS Precision Linear Polarizer is constructed by laminating a polymer polarizing film between two fused silica windows with high-efficiency broadband antireflection coating mounted in a 2-inch housing with a well-labeled transmission axis.

In order to obtain images at different polarization orientations, the polarizer is mounted in a COTS rotation stage from Newport. The RSP-2T Rotation Stage features a retaining ring to secure the optic. The polarizer can be coarsely aligned using the knurled edge of the rotating platform while fine adjustment is achieved with a precision adjustment knob. Angular position is indicated on a 360° scale graduated in 2° increments [1].

The full set of components was mounted using optical posts and post holders fastened to an optical rail and attached to a sturdy tripod for portability. Figure 3.7 shows the full assembly.

Figure 3.7:   Physical System Setup. A visible Sony CCD camera mounted behind a linear polarizer in a rotation stage.

3.3.2 Calibration .   Two types of calibration are required for this system. A radiometric calibration is done to change the digital counts given by the camera to a measurement of radiance in order to use the Stokes equations from Chapter II. A spatial calibration allows images to be rectified in order to use a pinhole model for multiple view geometry techniques.

3.3.2.1 Radiometric Calibration.   In order to calculate the Stokes vectors and subsequently the degree and angle of polarization, the measurements must be in the form of a radiance measurement. The Sony camera used for this polarimeter outputs 8-bit unsigned integer, digital count data. In order to convert these digital counts to radiance, a radiometric calibration must be performed.

To perform this calibration, the system was set up behind an integrating sphere with radiometric output known in terms of W/cm^2sr. The output of the integrating sphere was adjusted to 10 different levels and images were taken at each camera setting of interest. Figure 3.8 shows an image of the calibration setup.

There are many factors which can affect the digital output of the camera for a given radiance. Therefore, all camera settings were manually set and the compression techniques inherent in the digital system were included as part of the unknown system. Figure 3.9 shows the full image chain described by the calibration done in this section. A known radiance is input into the system in front of the polarizer, and digital counts are measured after processing.

In order to determine the radiance from a future image, a function is fit which corresponds the input and output of the system. Figure 3.10 shows the response of the camera as a function of shutter speed.

This figure shows that the response for this image chain is not a simple linear gain and offset, and that a more complex curve must be fit to the function. The log scale shown in Equation (3.1) was used for this purpose. This equation shows how the output from the camera in digital counts, D, is converted to radiance, R. The
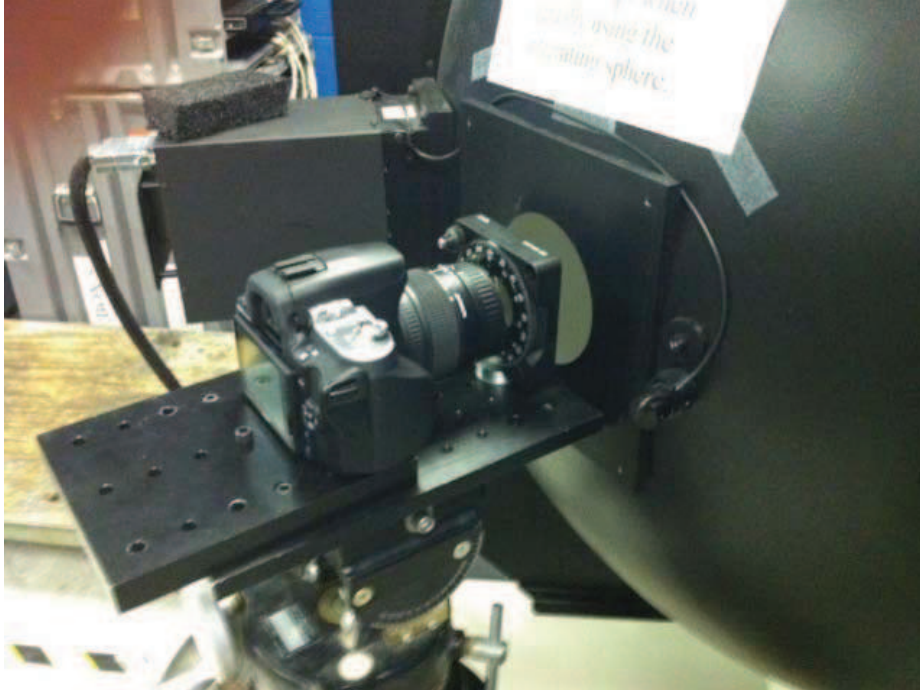
Figure 3.8: Spectral Calibration Setup. The camera system is placed behind an integrating sphere. Images are then taken at multiple settings.
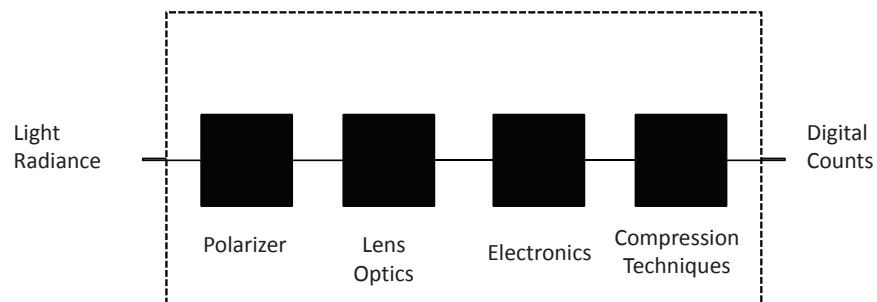
Figure 3.9: Physical Image Chain. Individual components are not calibrated in this paper. Instead the system is treated as a whole and a single calibration is done for one set of camera settings.
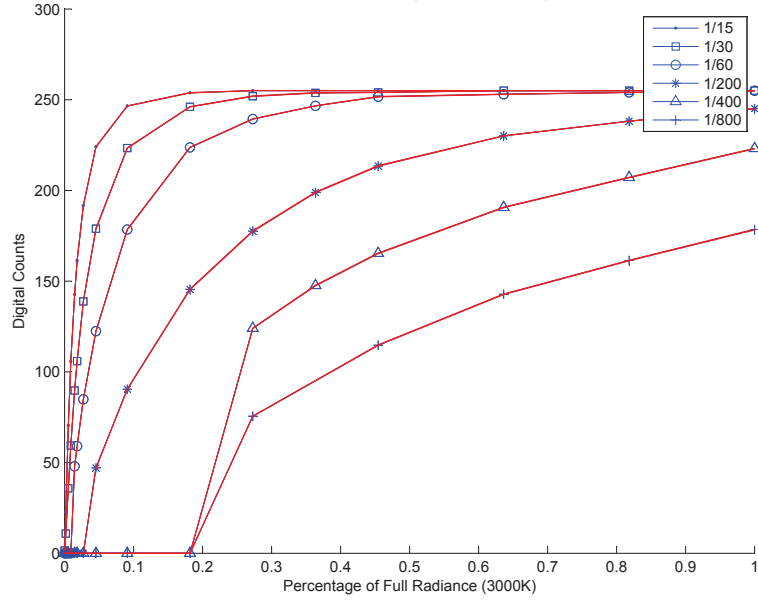
Figure 3.10: Camera Responses for given camera settings. These are the non-linear responses of the camera versus input radiance for different shutter speeds.

free parameters a and b are fit to each pixel for each camera setting to account for any differences due to camera parameters or compression techniques.

$$R = 10^{aD+b} \tag{3.1}$$

Each pixel of the camera has a slightly different response to radiance as it is processed through the image chain. Therefore, each pixel has its own calibration curve fit to it. Figure 3.11 shows the camera response for a subset of pixels with radiance fit equations overlaid on them. This figure shows a reasonable fit between the camera response and the log scale equation.

A concern over polarization imparted due to camera optics and camera noise was quickly laid to rest. Results of the radiometric calibration show that there is minimal additional polarization imparted by the camera. By finding the degree of polarization for a scene known to be completely unpolarized, the polarization imparted by the camera optics and noise can be determined. Figure 3.12 shows a DoLP image of

55
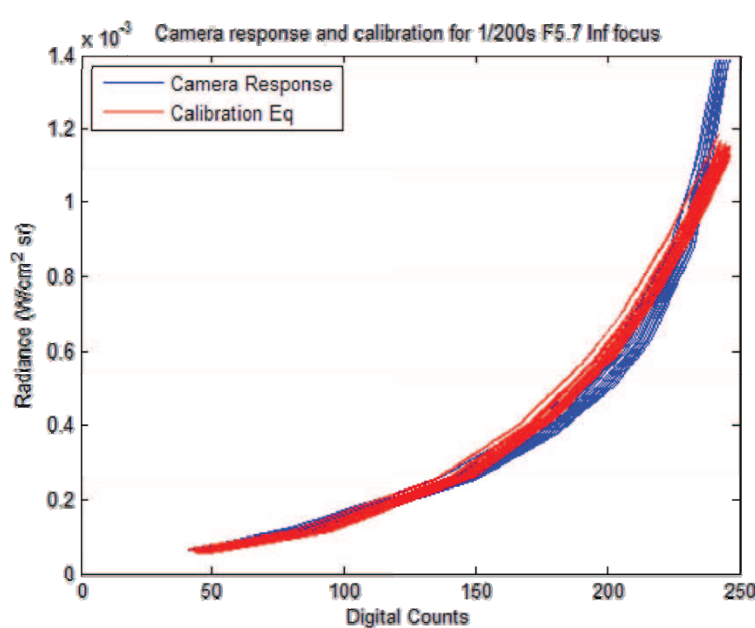
Figure 3.11:   Typical camera response and corresponding calibration equation. These calibrations are found using MATLAB's polyfit function.

the integrating sphere for one camera setting. There is a slight radial pattern in the image, however, the degree of polarization only ranges from $0 - 0.1\%$.

> *3.3.2.2   Spatial Calibration.*   The methods described in Chapter II for non-polarimetric shape recovery require a pin-hole camera model. Image distortion, due to optics, is removed through a spatial calibration. A calibration software package developed at CalTech was used for the spatial calibration and the results of that calibration are presented here.

Information on the CalTech calibration software can be found in [6]. The software requires a user to take images of a flat, checkered calibration board at a number of relative orientations. An image of the calibration board can be found in Figure 3.13. The software allows a user to input any number of images. It will then ask the user to find the four corners of the checkerboard pattern. A corner detection algorithm is performed by the software to find the corners within the user-defined square. Differences in the actual corners and the interpolation from the extreme corners relate
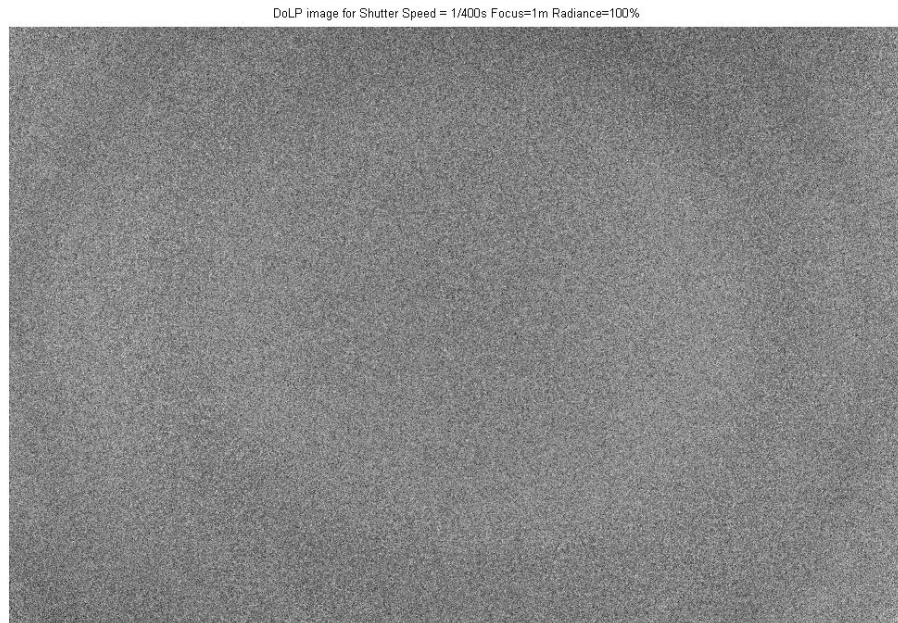
Figure 3.12: Degree of polarization measurement of the integrating sphere. This figure shows that there is minimal polarization imparted from the optics, electronics and compression from the imaging system. The degree of polarization in this image ranges from $0 - 0.1\%$
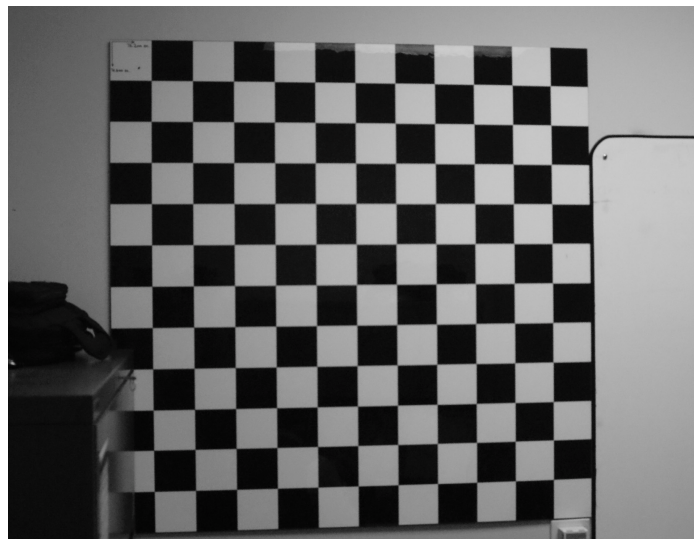
Figure 3.13: Camera Calibration Board. Imaging this board from multiple views allows the Caltech Calibration software to determine intrinsic camera parameters that can be used to remove image distortion.

directly to the distortion parameters of the lens. The software calculates the best set of distortion parameters which fit all given images and outputs a best estimate of the camera parameters and the errors associated with these estimates.

The CalTech software outputs four main intrinsic parameters for the camera. The focal length is given, in terms of pixel size, for the x and y direction. The principal point defines the center of the focal plane given in pixel number. The skew defines the angle between the x and y orientations of the focal plane. A skew of 0 relates to a $90°$ angle between these axes. The distortion parameters are a set of five coefficients used to define the radial and tangential distortion.

A calibration is only good for a particular set of lens parameters. The zoom, focus, and aperture all affect these characteristics and must be noted carefully. In order to determine any distortion effects imparted by the linear polarizer, a calibration was done using the same set of camera locations at each of the four polarizer orientations. Results of these calibrations can be found in Table 3.1. These results show that there are no differences in distortion as a function of the angle of the linear polarizer. This means that a single spatial calibration can be used for all four orientation images as well as the degree and angle of polarization images.

3.3.3 Processing and Products .   The image processing chain is accomplished through a spectral calibration, computation of the Stokes images and calculation of the degree and angle of polarization. Once the radiometric calibration is done for a particular set of camera parameters, images of calibration coefficients are stored. These images are then processed through the radiometric calibration using Equation (3.1). The calculation of the Stokes images are done by incorporating the relationships of the 0, 45, 90 and $135°$ orientation images found in Section 2.4.2. Finally, the degree and angle of polarization images can be found from the Stokes images by use of Equations (2.24) and (2.25). Examples of these products are found in Figure 3.14.

In general, these images work for computer based algorithms. However, they are sometimes difficult for a human to interpret. In order to advance human un-

Table 3.1: Spatial Calibration Results. These parameters show that here is little difference in the spatial calibration due to the angle of the linear polarizer. This allows for a single calibration to be used for each image.

Orientation	Focal Length $\begin{pmatrix} x \\ y \end{pmatrix}$	Principal Point $\begin{pmatrix} x \\ y \end{pmatrix}$	Distortion
0°	1154.20 ± 1.83	386.91 ± 2.39	-0.11372 ± 0.00831
	1156.66 ± 1.73	256.57 ± 2.69	0.08415 ± 0.07144
			-0.00135 ± 0.00052
			0.00035 ± 0.00048
			0.00000 ± 0.00000
45°	1155.14 ± 1.83	386.52 ± 2.39	-0.11230 ± 0.00830
	1157.48 ± 1.72	255.30 ± 2.69	0.08152 ± 0.07152
			-0.00135 ± 0.00053
			0.00038 ± 0.00048
			0.00000 ± 0.00000
90°	1154.81 ± 1.81	385.16 ± 2.36	-0.11280 ± 0.00818
	1157.13 ± 1.70	255.06 ± 2.65	0.07480 ± 0.07045
			-0.00143 ± 0.00052
			0.00013 ± 0.00047
			0.00000 ± 0.00000
135°	1154.87 ± 1.81	385.22 ± 2.36	-0.11415 ± 0.00818
	1157.22 ± 1.70	254.12 ± 2.65	0.08920 ± 0.07047
			-0.00169 ± 0.00052
			0.00019 ± 0.00047
			0.00000 ± 0.00000
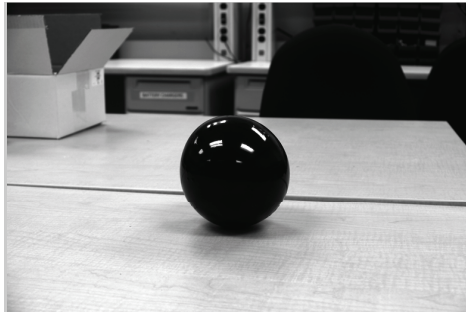
(a) Intensity

(b) Degree of Polarization
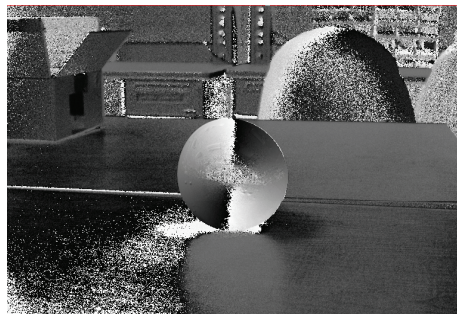
(c) Angle of Polarization

Figure 3.14:   Polarization Products Examples. These images show the intensity, degree and angle of polarization for a glossy black painted sphere on a wooden table in a well lit room. The degree of polarization increases towards the edges of the sphere and the angle of polarization changes gradually around the sphere from 0° to 180°, both as expected.

Figure 3.15: Hue, Intensity, Saturation Pseudo-color representation of the black ball shown in Figure 3.15.

derstanding in a single image a pseudo-color interpretation was developed [5]. This image uses a hue, intensity, saturation color map to produce an image, which shows the intensity, degree and angle of polarization in a single image. The hue of the image is related to the angle of polarization. The intensity is simply the Stokes S_0, or intensity image. The saturation is related to the degree of polarization. Therefore, images with a deeper color convey a higher degree of polarization.

Figure 3.15 shows an example of this color product applied to the glossy black ball shown in Figure 3.14. This image shows the same increase in degree of polarization toward the edges of the sphere and continuous rotation of the angle of polarization. In order to show color overlaid on the black sphere, the intensity image has been reversed so that dark objects show up lighter and lighter objects look darker.

The tools presented in this chapter describe the foundational tools needed for an efficient research effort. Simulations provide an easy medium to test algorithms, while the physical system allows for more complex environments and anomalies not covered in simulation. The next chapter describes the tests performed with these tools. It

61

shows how the ultimate goal of attitude estimation is achieved through understanding of multiple steps and simpler problems.

IV. Methodology of Tests

An iterative approach is taken to reach the goal of using polarimetric measurements to achieve attitude estimation. The first set of tests are performed to determine how well Shell model surface parameters can be estimated given limited information about the test subject (Section 4.1). It was determined that without a complete set of surface parameters, the only useful surfaces for simple reconstruction and navigation would be flat surfaces. An algorithm is developed to find flat surfaces in an image (Section 4.2). A set of constraining assumptions is used to simplify the Shell model and determine surface orientation (Section 4.3). Finally, assumptions are made about an indoor environment which allowed for camera orientation to be estimated (Section 4.4). Each previous test proves useful in providing vital information for the next test. This Chapter describes the methodology used for each of these tests. Results for each test will be discussed in Chapter V.

4.1   Estimation of Intrinsic Surface Parameters

In order to perform the ultimate goal of estimating receiver geometry, some information must be known about the intrinsic surface parameters. A set of MATLAB GUIs were developed to test the limits of estimation of the surface parameters, and to determine how much must be known about a situation prior to estimating the surface geometry.

The simulations started slowly, with a GUI developed to determine estimation of a single surface parameter given knowledge of the other parameters and the camera-surface-source geometry. More GUIs were built in succession to try to estimate multiple parameters, to use multiple measurements in the estimation, and to try a multiple hypothesis testing algorithm.

The estimation algorithm used for these tests is presented in Section 4.1.1. Each MATLAB GUI used for parameter estimation is described in detail in Section 4.1.2, and results of these tests can be found in Chapter V Section 5.1.

4.1.1   Algorithm .   Surface parameters are estimated through the Levenburg-Marquardt method [22]. In simple terms, the Levernburg-Marquardt algorithm is a non-linear minimization algorithm. The algorithm used here determines a Jacobian numerically. It then uses this information to determine a state change which most efficiently converges to a minimum.

This algorithm is used to reduce the error in measurements between actual and estimated values for surface parameters or geometry. Given a set of 'known' parameters, an initial guess at the parameters of estimation can be processed through the non-linear set of equations in the Shell target model. The output of these equations is subtracted from the given measurement, or the measurement determined from the actual parameter set. The difference in measurements, known as the residual, is the parameter to be minimized by the Levenburgh-Marquardt algorithm.

Figure 4.1 shows the error in the degree of polarization measurement as a function of error in estimation of the real part of the index of refraction, n. A linearization is shown at the starting estimate for the index of refraction and illustrates how the algorithm eventually converges to a minimum in the measurement error.

4.1.2   MATLAB GUIs .   Three distinct GUIs were developed to determine the limits of estimation of intrinsic surface parameters. The first GUI shows how errors in a single parameter relate to errors in individual measurements given the rest of the parameters and the surface geometry are known (Section 4.1.2.1). The same GUI allows a user to show the final error in estimation as a function of the actual parameter and the starting estimate. The second GUI allows a user to choose the set of known and unknown parameters in order to estimate multiple parameters (Section 4.1.2.2). The final GUI uses the concept of multiple hypothesis testing to determine a best fit set of parameters and therefore a best material, given actual and estimated conditions (Section 4.1.2.3).
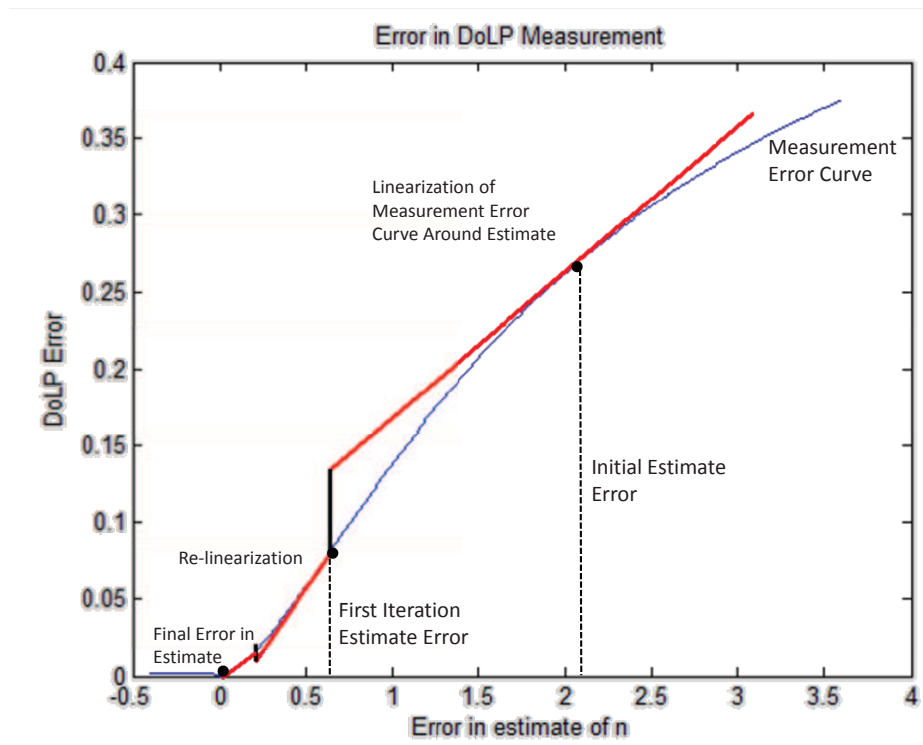
Figure 4.1:   Illustration of the Levenburgh-Marquardt Algorithm. The linearization and residual determine which way and how far to propagate the estimation. This is done iteratively until the algorithm finds a local minimum.

4.1.2.1 Single Parameter Estimation . There are eight parameters in the Shell Target Model that are used to represent the surface parameters. These parameters were described in detail in Chapter II, Section 2.5. This GUI was developed to determine how estimation of a single parameter is affected by other parameters. It also gives a user a better understanding of how each parameter affects the intensity, degree of polarization and angle of polarization. The GUI shown in Figure 4.2 also allows a user to determine the errors between measurements using actual surface parameters and estimated surface parameters.

This GUI can be used to determine constraints that can be used to better estimate a given parameter. Using an estimation technique to iteratively determine a final solution of a surface parameter given an initial estimate relies on a surface plot like the one shown in Figure 4.2. Valleys in this plot are where an estimate would eventually converge. Starting with an estimate on the wrong side of a peak in this image would lead to a wrong final estimation.

A non-linear least squares algorithm such as Levenburg-Marquardt can be used to try to estimate all or some of these parameters. However, these techniques require a relatively 'close' starting estimate. An option within this GUI was used to determine how far off a starting estimate for a single parameter could be and still converge to the correct solution. This option, shown in Figure 4.3, allows the user to set the geometry and each of the 'known' surface parameters. It then varies a single surface parameter and a starting estimate for that parameter and outputs a graph displaying the error in the final estimate. The limit of starting estimates can easily be seen.

4.1.2.2 Multiple Parameter Estimation . Once a user has an understanding of how a single parameter can affect the measurements, the GUI shown in Figure 4.4 can be used to determine the extent to which multiple parameters may be estimated. It also allows for a user to determine observability of parameters under certain conditions and allows for determination of a viable set of constraints in order
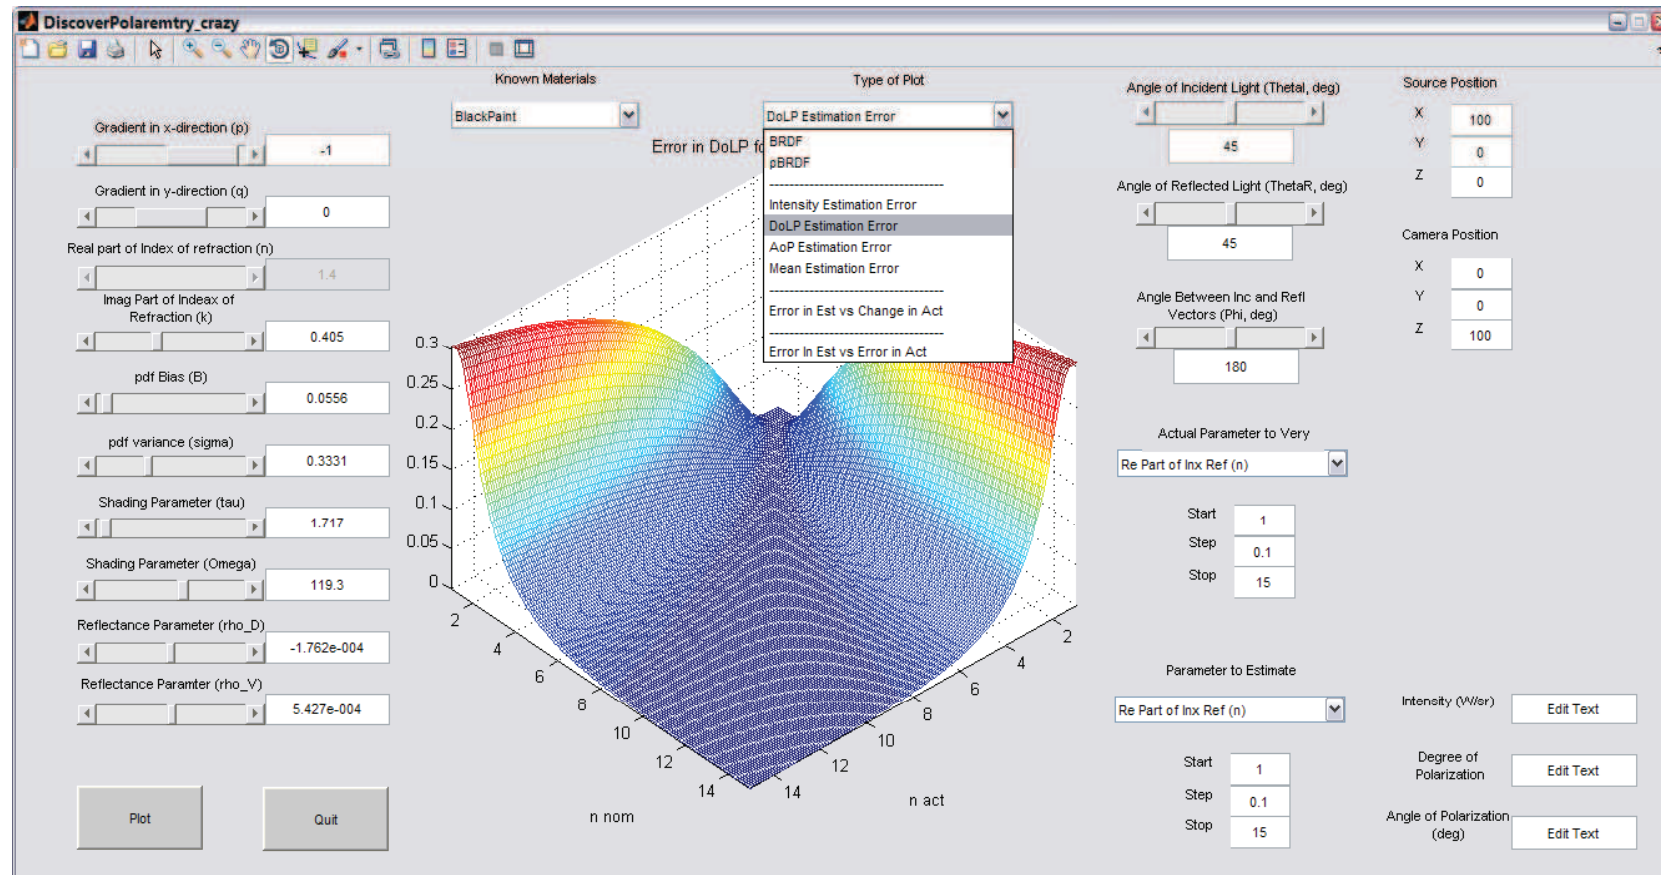
Figure 4.2:   MATLAB GUI to give a user a feel for how parameters affect the intensity, degree and angle of polarization measurements. This GUI also allows a user to determine viable constraints for estimation.
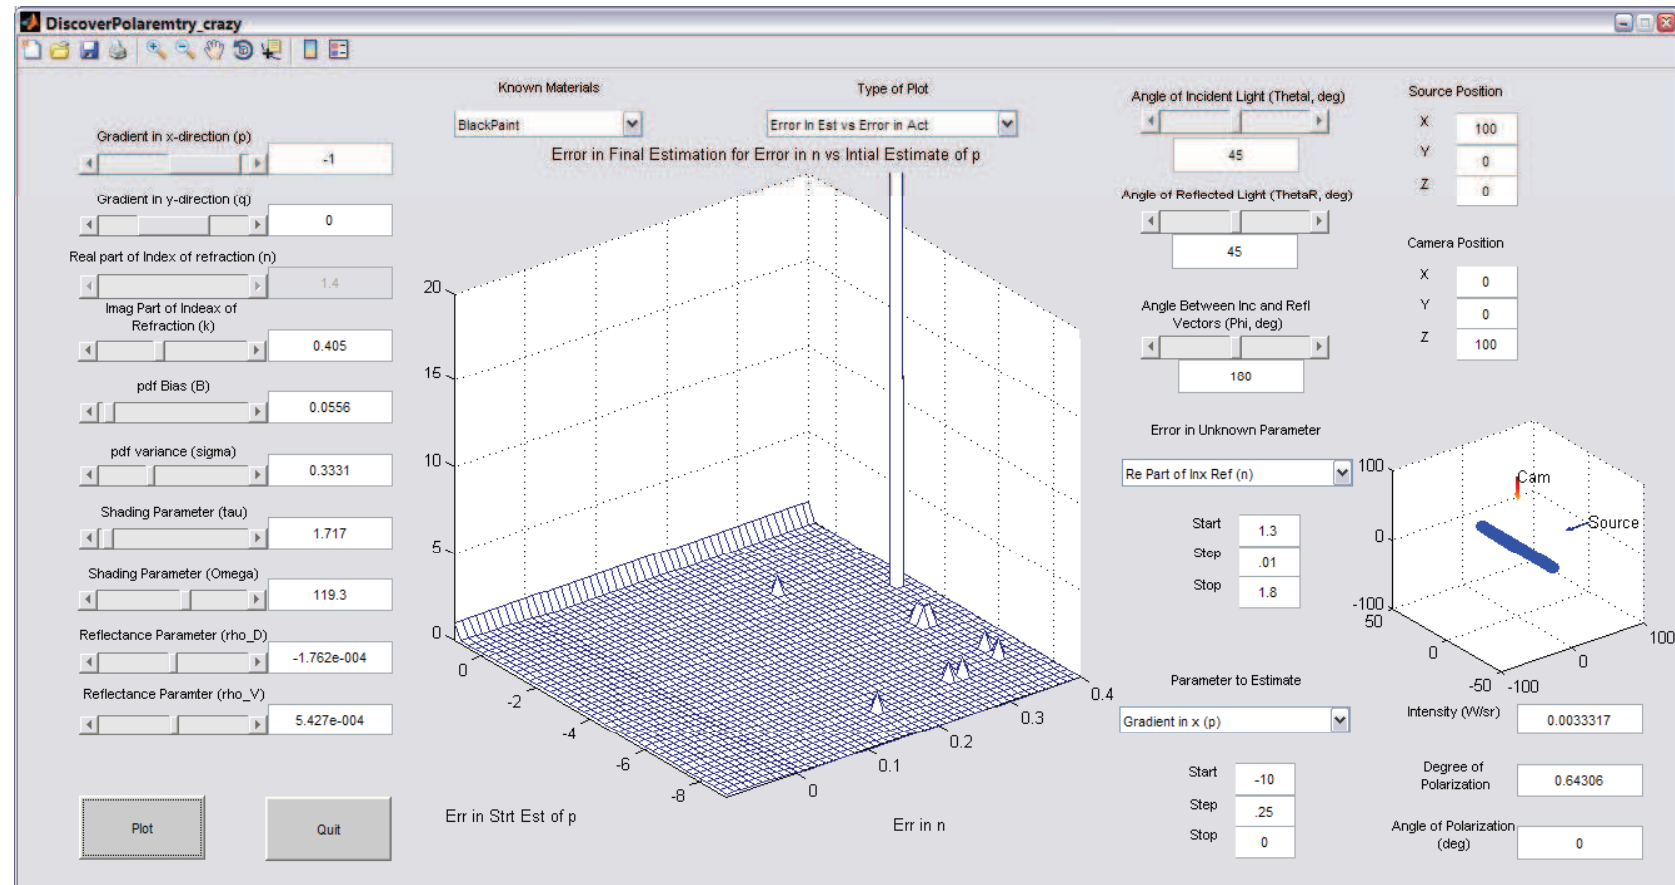
Figure 4.3: Single Parameter Estimation GUI esample. This GUI shows a user the extent of starting estimates that converge to a true value for a given parameter.

to estimate the desired parameters. The estimation algorithm used for the GUI is the same as estimating parameters in Section 4.1.

This GUI allows a user to define an actual set of surface parameters and geometry and an estimate of 'known' and unknown parameters. It allows the user to choose the parameters to estimate and even allows for multiple measurements at user defined geometries. Using this GUI, it is easy to determine which parameters are most important in the estimation of other parameters or geometry. This allows a user to determine constraints that can be placed on Shell parameters such as the material makeup, smoothness, reflectance or shadowing of a surface.

4.1.2.3 Multiple Hypothesis Testing Method . The ultimate goal of this thesis is to determine camera orientation using what little is known about the scene. It was shown using the GUI from Section 4.1.2.2 that the best results for geometry estimation are given when a full set of surface parameters are known. Given that a limited number of known materials can be found in an indoor environment, a multiple hypothesis test can be performed to determine a full set of surface parameters.

The GUI developed in this section, shown in Figure 4.5, allows a user to define actual surface parameter values and a 'known' geometry. It will then calculate errors in actual and estimated measurements and chose the material that best fits the measurements.

Even given a full set of Shell target parameters, the above GUIs show that there is some error in geometry estimation for certain orientation and lighting conditions. In order to simplify the estimation of orientation of a surface, only large flat surfaces were considered. The next section describes the techniques and tests done to identify flat surface in an image.

4.2   Determining if a Surface is Flat

It was shown in Chapter II Section 2.3 that there are already techniques to determine relative orientation of sequential images and surface geometries if a surface
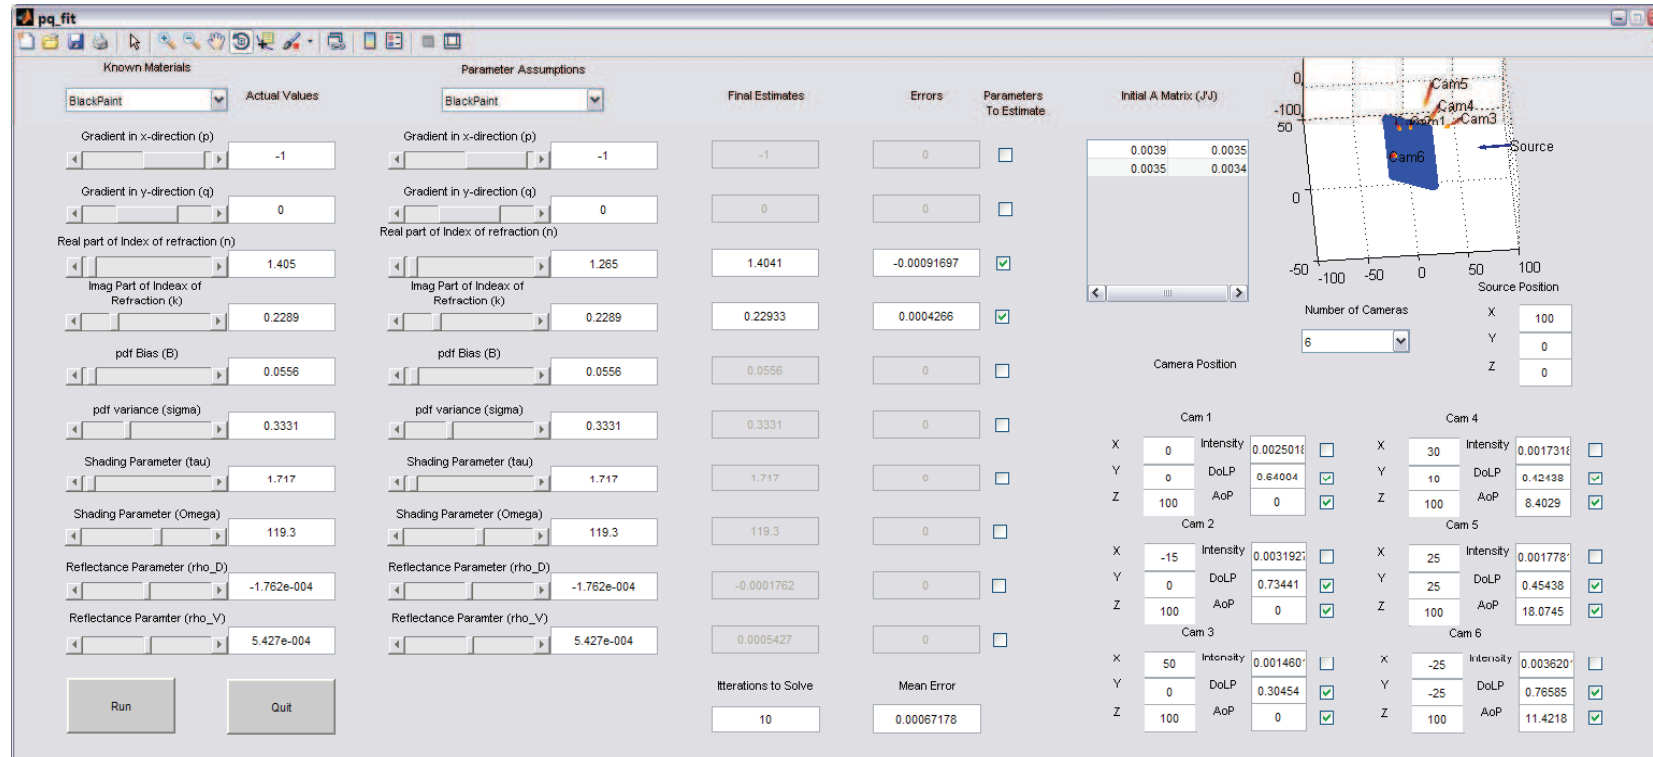
Figure 4.4:   MATLAB GUI to Estimate Multiple Parameters. This GUI is used to estimate a user defined set of parameters given knowledge of the set of known parameters. This GUI allows a user to determine observability of parameters and limits of viable constraints in order to create better estimation techniques.
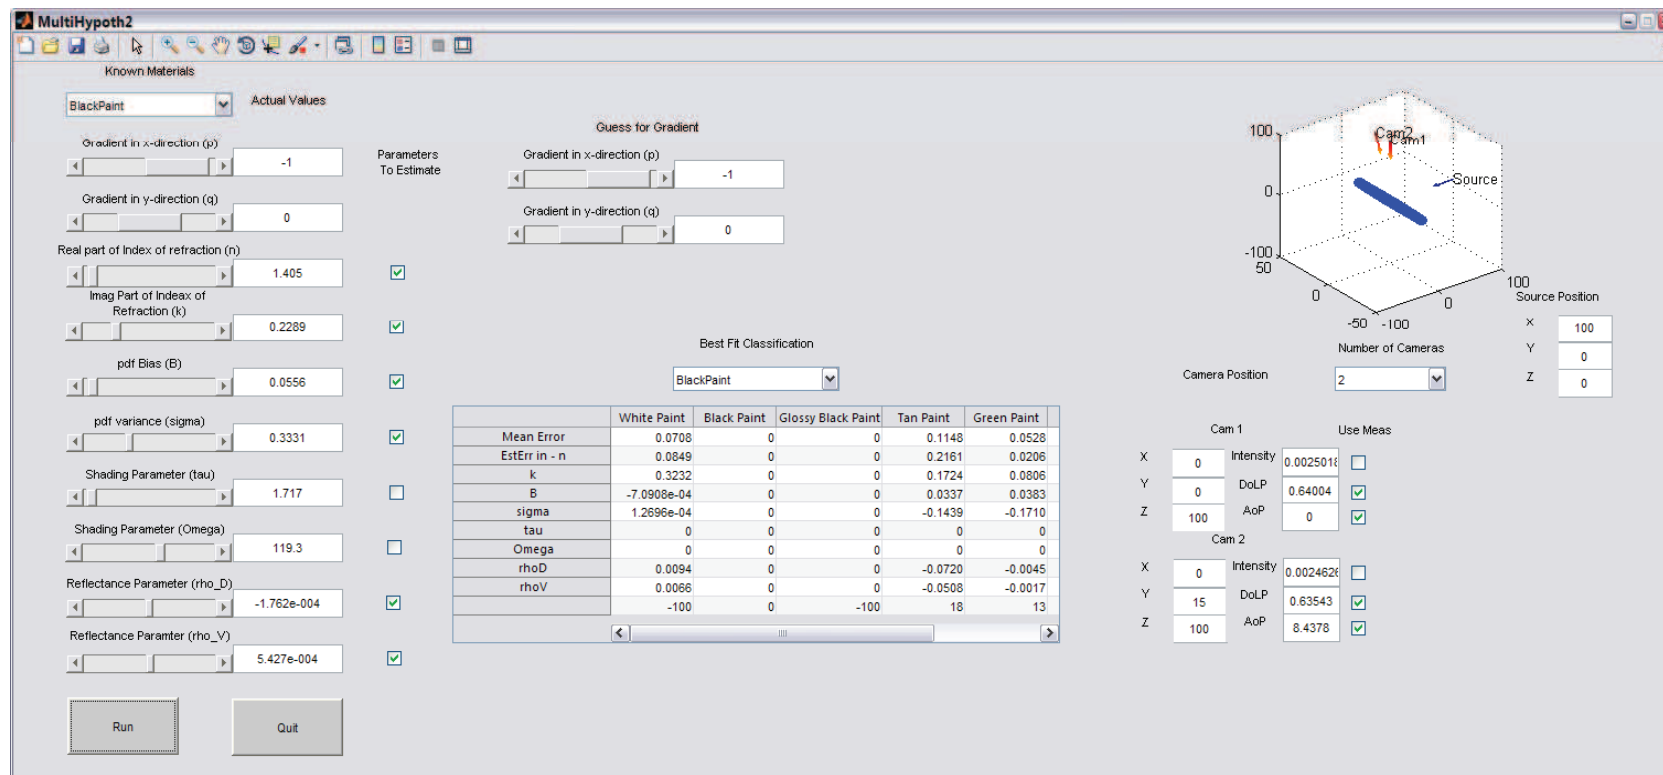
Figure 4.5:   Multiple Hypothesis GUI Example. This GUI allows a user to input actual parameters and geometry for a surface. It will calculate the expected measurement for a given set of targets known to be in the scene and return the target material that most closely fits the given measurement.

can be identified as flat. It is also easier to determine surface orientation using these large flat structures, and they happen to be prevalent in man-made environments. The goal of this section is to determine which parts of an image contains large flat structures.

Two methods of determining flat surfaces were employed. The first uses a windowing technique in which windows with a standard deviation for degree and angle of polarization below a set threshold are considered to be flat. The second method uses a gradient of the degree and angle of polarization images. Any pixel with a combined weighted gradient below a given threshold is considered to be flat. Pixels are then binned by angle of polarization and grouped by connected components.

With the windowing method, for illustration, arrows are placed at the center of windows that are determined to be flat. The orientation of the window points in the direction of the surface normal, projected into the focal plane. The orientation is given by the angle of polarization, and the length is proportional to the degree of polarization.

This method was tested with the physical polarimeter on two glossy black objects, a flat, painted plate and a painted ball. As a demonstration, a point was manually selected on each image and surface flatness was determined. Figure 4.6 shows the points chosen on each surface. The flat plate shows an arrow in the direction of the surface normal because it was determined to be flat. The ball only shows an outline of the window, without an arrow, because it was determined to be curved.

These tests were then expanded to a more complex indoor environment. Figure 4.7 shows an image of hallway that was used with the windowing algorithm. It can be seen immediately that not all surfaces that are flat are determined to be flat. The hallway does not present a complete set of specular geometries. However, for areas with specular reflections, it does find flat surfaces. It also does not present any false positives. That is, it does not determine that any location is flat if it is in fact curved.
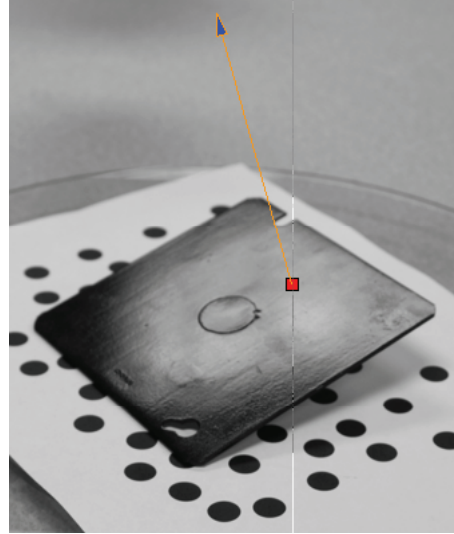
(a) Glossy Black Ball　　　　　　(b) Glossy Black Metal Plate

Figure 4.6:　　Glossy Black Ball and Plate in Flatness Test. These images show how the windowing test performed against a ball and a flat plate. The red box on the ball shows that it was determined to be not flat, while the arrow on the flat plate shows that it was determined to be flat with a surface normal in the direction of the arrow.

The second method of determining if a surface is flat, the gradient test, is generally easier for a human to interpret. Figure 4.8 shows the same hallway having gone through the gradient flatness algorithm. Regions of connected flat components with similar angles of polarization are grouped and colored. Arrows are placed at the center of these groups and point towards the surface normal of the group.

The grouped flatness test tended to be more reasonable for use in determining surface or camera orientation. The smaller number of flat areas and larger number of pixels per area mean that there are fewer and less noisy measurements.

4.3　Estimation of Surface Orientation

The goal of the following tests is to explore the limits of estimation of pitch and tilt angle. These limits along with the assumptions presented in Section 4.4

Figure 4.7: Typical hallway results for the standard deviation of a sliding window test. The arrows show the places that were determined to be locally flat and are pointed in the direction of the surface normal. Arrows in the image may be pointing in a direction opposite to the surface normal due to the ambiguity of the angle of polarization.
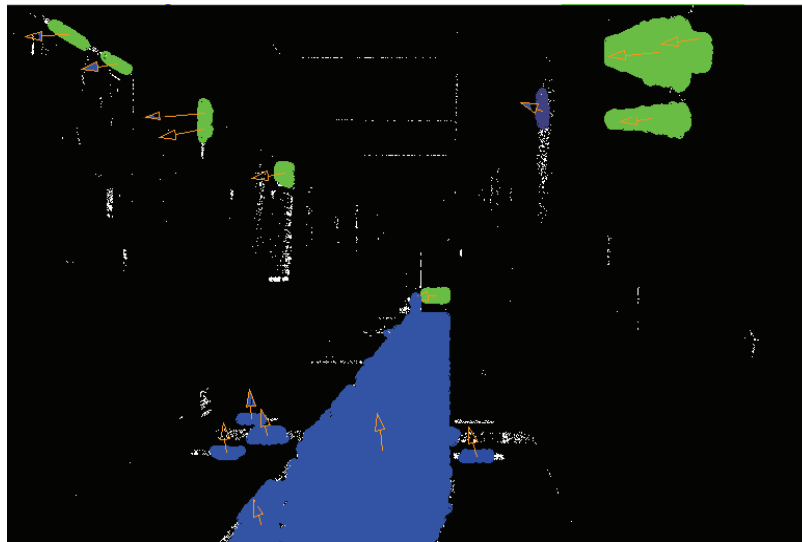
Figure 4.8: Typical hallway results for the gradient of degree and angle of polarization threshold. Connected components with similar angle of polarization are grouped. The arrows are at placed at the center of the groups and points towards the center of the groups surface normal. Arrows in the image may be pointing in a direction opposite to the surface normal due to the ambiguity of the angle of polarization.

74

will ultimately be used to determine camera orientation and the uncertainty in the estimation of that orientation.

Numerous tests were performed using all of the tools described in Chapter III. The MATLAB GUIs created for these tests are described in Section 4.3.2.1. These GUIs are verified using DIRSIG simulation software in Section 4.3.2.3. Finally, the physical polarimeter was used to determine any inconsistencies between simulation and reality (Section 4.3.2.4).

The preferred method of representing surface orientation up to this point was by surface gradient in the x and y directions of the camera frame of reference. This allowed a single global frame of reference for the surface, source and multiple cameras. However, for a single view with specular geometry, it is easier to describe the relative orientation between the surface and source in terms of a pitch and tilt angle. Instead of estimating the surface gradient in the camera Cartesian coordinates, as in Section 4.1, it is beneficial to think of the surface geometry in terms of pitch angle, or the angle between the focal plane of the camera and the surface normal, and tilt angle, or the angle between the zero degree angle of polarization and the surface normal. The angle of polarization describes the angle of the in-plane polarization state, and is therefore directly related to the surface tilt. Estimation of the surface pitch is more complicated and is the main contribution of this section.

Two algorithms are presented in Section 4.3.1 to describe the estimation of pitch angle. An algorithm to be used if a full set of Shell parameters are available, and a simplified method if the full set of parameters are unknown but the surface is known to fall within a set of constraints generally found in indoor environments. These algorithms are then tested using each of the tools presented in Chapter II.

4.3.1 Algorithm for Model Simplification .   If a full set of Shell target parameters are known and a specular geometry exists, the Levenburg-Marquardt estimation techniques from Section 4.1 were shown to provide good results in the estimation of pitch angle. The equation used to predict degree of polarization, used in the

75

Levenburg-Marquardt algorithm, was changed to allow an input of pitch angle, as opposed to surface gradient and source and receiver pointing angles, as in Section 4.1.

This algorithm only works if a source is considered to be in a specular direction. This is a reasonable assumption considering that for very smooth or glossy materials the vast majority of polarization is only presented in a specular direction and off-specular geometry will normally yield little polarimetric information. Given that there is a light source in the specular direction, the source-surface-camera geometry is such that $\theta_i = \theta_r = \beta$ and $\Delta\phi = 180°$, and all other geometric parameters can be neglected.

It was shown using the GUI in Section 4.1 that if a full set of Shell parameters are not known well, that errors in guessing these parameters will quickly lead to errors in estimation of geometry. If the full set of Shell Target model surface parameters are unknown, but the surface is known to be composed of a smooth dielectric, the Shell model can be reduced back to the Fresnel Reflectance Model and assumptions can be made about the complex index of refraction.

If the material is smooth relative to the wavelength of light used, the surface roughness and shadowing parameters of the Shell model can be reasonably neglected and the polarization component of the Shell Target model can be simplified back to the Fresnel Reflectance equations. This reduces the parameters required to describe the degree of polarization to only three, the angle of reflection, β, and the two components of the complex index of refraction, $n + ik$. If a dielectric surface can be assumed, the values of n and k can be constrained to $n \approx 1.5$ and $k \approx 0$.

Given these assumptions, the Fresnel reflectance equations, shown in Equations (2.27) - (2.28), reduce to

$$r_s = \sqrt{\frac{(A - \cos(\beta))^2 + B^2}{(A + \cos(\beta))^2 + B^2}} \qquad (4.1)$$

76

$$r_p = \sqrt{\rho_s \frac{(A - \sin(\beta)\tan(\beta))^2 + B^2}{(A + \sin(\beta)\tan(\beta))^2 + B^2}} \tag{4.2}$$

where A, B, C and D are functions of the complex index of refraction and are given by

$$\begin{aligned} A &= \sqrt{(C+D)/2} \\ B &= \sqrt{(C-D)/2} \\ C &= \sqrt{4n^2 k^2 + D^2} \\ D &= n^2 - k^2 - \sin(\beta)^2 \end{aligned} \tag{4.3}$$

The Mueller matrix, which describes how the Stokes vector changes through a reflection is then reduced to

$$M = \begin{bmatrix} \frac{r_s^2 + r_p^2}{2} & \frac{r_s^2 - r_p^2}{2} & 0 \\ \frac{r_s^2 - r_p^2}{2} & \frac{r_s^2 + r_p^2}{2} & 0 \\ 0 & 0 & r_s r_p \end{bmatrix} \tag{4.4}$$

4.3.2 Tests . Each of the tools described in Chapter II was used in the testing of this algorithm. Several MATLAB GUIs were developed to quickly determine how well the estimation of β could be performed. These GUIs were validated using a DIRSIG simulation and then the physical polarimeter was used to determine if there were any differences between the simulation and reality.

4.3.2.1 MATLAB GUIs . Two sets of GUIs were designed and built to show how estimation of β angle changes as a function of model parameters. The first GUI was developed to compare differences between DoLP for a material with an actual set of Shell Target parameters and an estimate of these parameters. Figure 4.9 shows a representation of this GUI.
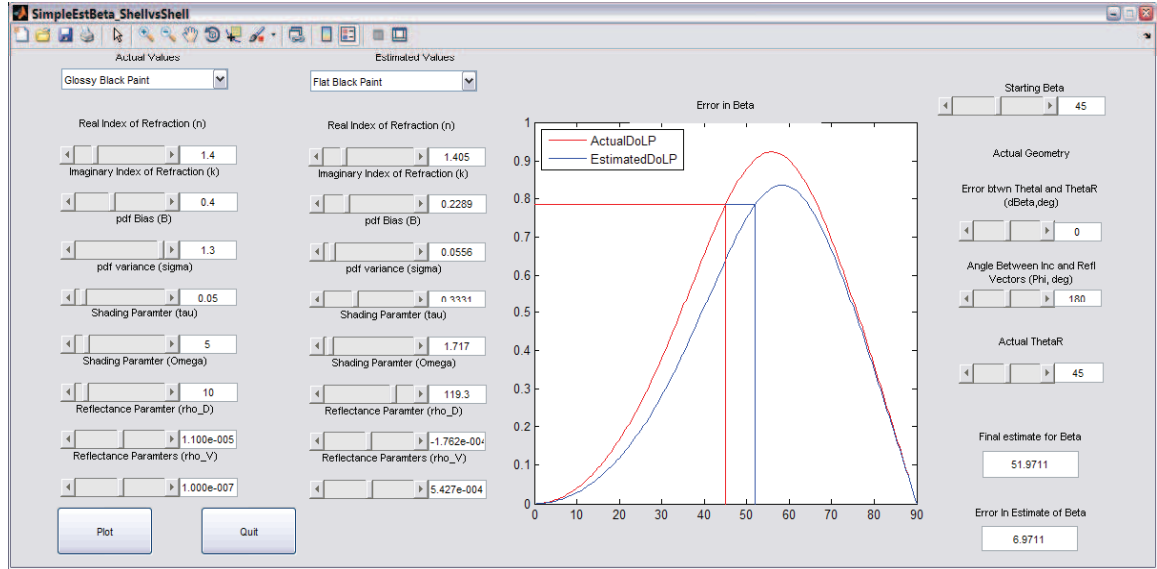
Figure 4.9:   MATLAB GUI to Estimate β Using a Set of Actual and Estimated Shell Parameters. This GUI allows a user to input a set of actual and 'known' shell parameters and determine how error in estimation of the actual parameters related to error in estimation of β.

If a full set of Shell parameters are not known, but if the assumptions of a specular geometry and a smooth specular surface are upheld, there are still two surface factors that are not generally precisely known. The complex index of refraction of a surface contains the only two factors left unknown before estimating β.

The MATLAB GUI shown in Figure 4.10 shows a simulation design that allows a user to input Shell Target parameters and geometry for a surface and starting estimations for the complex index of refraction of the surface and the angle of reflection off the surface. The figure within the GUI shows the actual degree of polarization as a function of θ_r and the estimated degree of polarization as a function of β using the Fresnel Reflectance Equations.

This application allows the user to input actual and estimated values for the complex index of refraction. The user can then visualize the differences in DoLP as a function of β. It can also be seen that the initial guess for β is an important factor in this estimation. Choosing a β on the wrong side of the angle of maximum polarization, Brewster's angle, will yield an incorrect final result.

78
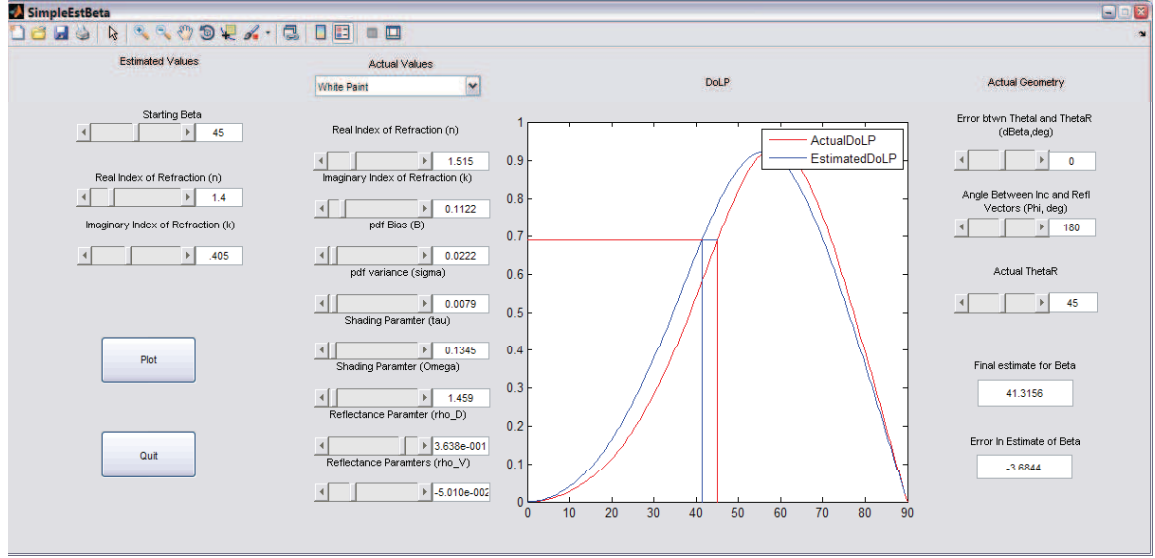
Figure 4.10: GUI to Show Estimation of β. This GUI estimates β angles given the actual Shell Target Model parameters and geometry and estimates of the complex index of refraction.

When the assumptions in Section 4.3.1 are violated, errors will result in the estimation of the surface normal. Two more GUIs were created to determine the extent of errors in the estimation of β given either an estimate of the full set of Shell Target parameters or a set of Fresnel parameters and an uncertainty in the actual set of Shell parameters.

4.3.2.2 Envelope of Errors in Estimation of Beta . Given an uncertainty in the actual Shell target parameters, there is some uncertainty in the error of the estimation of β. These uncertainties in error will eventually be used to determine measurement error strength when using β to estimate camera orientation in Section 4.4.

The GUIs shown in Figures 4.11 and 4.12 allow a user to define an actual set of Shell surface parameters and surface geometry. The GUI in Figure 4.11 will then vary the actual Shell target parameters and estimate a β angle using the simplified Fresnel method. It does this for 500 trials at each 1° interval of β and plots the mean and standard deviation of error in the estimate.

79
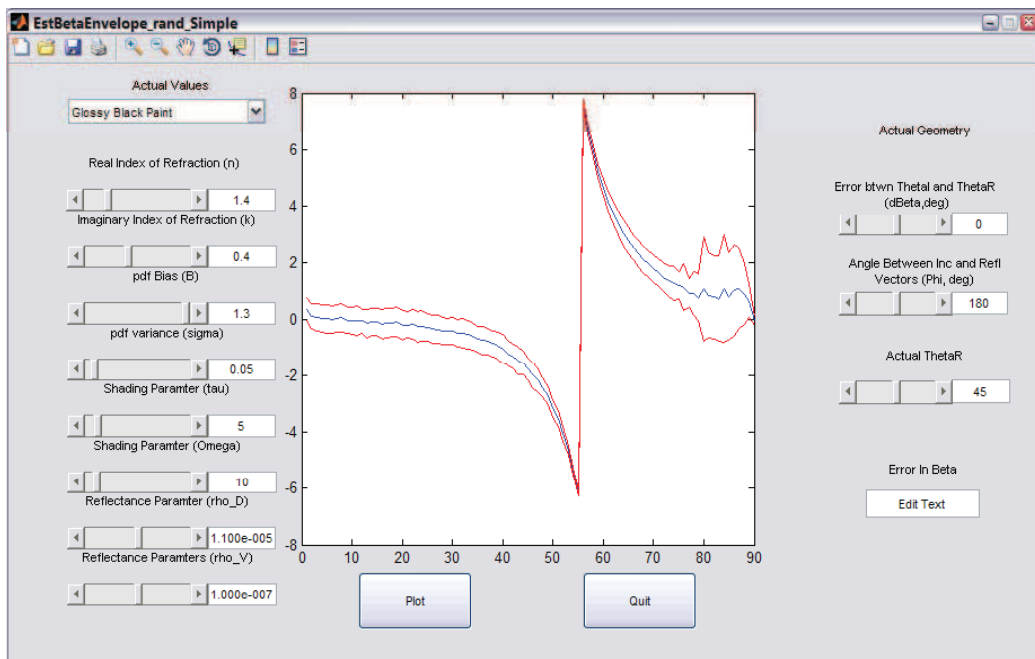
Figure 4.11:   Envelope of the Estimation of β Angle Given Fresnel Model. This GUI is used to estimate the mean and standard deviation of errors in the estimate of β given the Fresnel estimation method. The center line represents the mean of the error in β estimation, while the surrounding lines represent the mean plus a 1σ standard deviation.
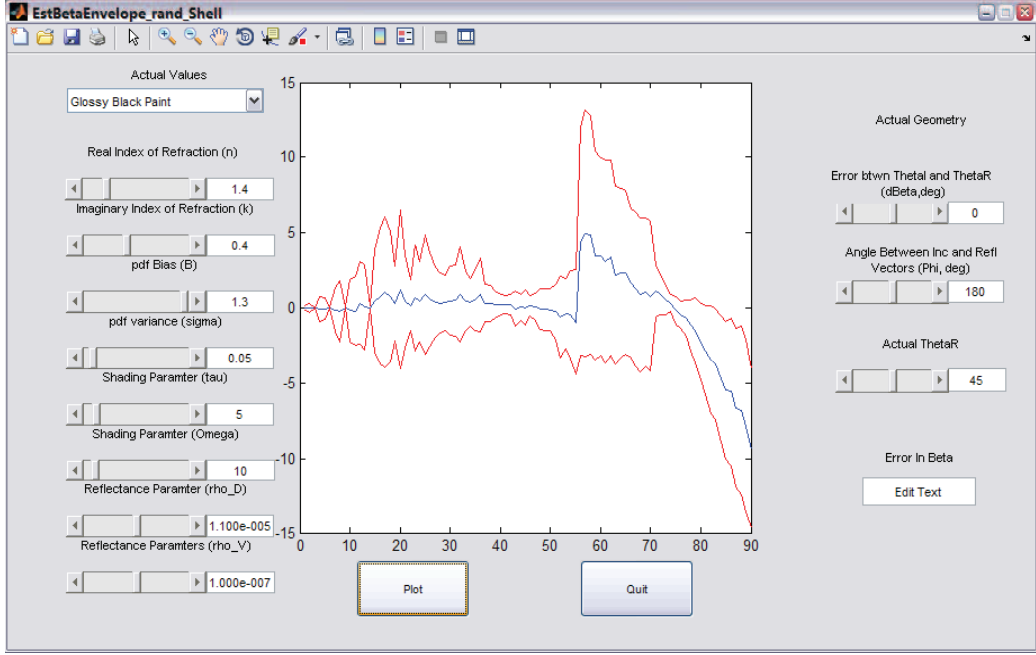
Figure 4.12: Envelope of the Estimation of β Given the Shell Model. This GUI estimates β angles given an estimate of the full set of Shell parameters. The center line represents the mean of the error in β estimation, while the surrounding lines represent the mean plus a 1σ standard deviation.

The GUI in Figure 4.12 gives the user the same input controls and assumes that the estimates for the Shell parameters are close to the actual parameters. It will vary the actual parameters and estimate β angle using the full Shell target model estimation technique. It then creates a plot of mean and standard deviation of error in these estimates, similar to the previous GUI.

A set of simulations and physical experiments were performed to prove the results found in Section 4.3.2.1. The simulation was performed in DIRSIG using a set of glossy black objects and positioning the camera and sun at various $\Delta\phi$ and β angles. An explanation of the DIRSIG implementation is presented in Section 4.3.2.3. An explanation of the experiments performed with the physical polarimeter are shown in Section 4.3.2.4.

4.3.2.3 DIRSIG Simulation . The simulations for this experiment show a set of three glossy black objects: a cube, a cylinder and a sphere. The
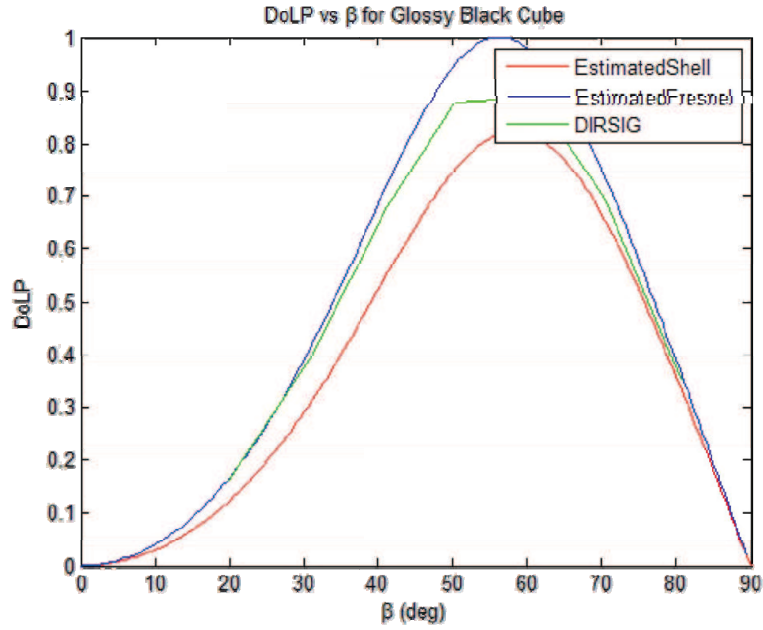
Figure 4.13: DoLP vs β for the DIRSIG Simulation of Fresnel and Shell Models. The DIRSIG simulation results show more agreement with the Fresnel model than with the Shell model with perfect geometry.

geometry of the system is set up at seven different β angles between $20 - 80°$. The source and camera are positioned at a specular orientation relative to the tops of the cube and cylinder.

For each of these images, the average DoLP of the top of the cube was taken and plotted versus the β angle. Figure 4.13 shows this plot along with the plots from the Fresnel and Shell estimates for a perfect geometry. It can be seen that the DIRSIG estimation, which does not have perfect geometry, still falls between the two models.

Three-dimensional surface normals were then found at user-defined positions. The surface normal can be found in the camera reference frame by using the angle of polarization to get the x and y vector components, and the estimation of β to get the z component. Figure 4.14 shows a representation of the surface normal on top of the cube. Since camera location and orientation are known perfectly in DIRSIG, a transformation of the surface normal can be done to get the normal in the local
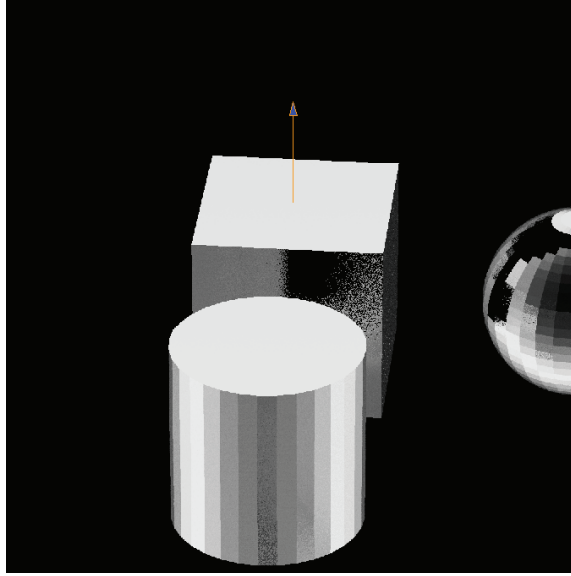
Figure 4.14: DIRSIG DoLP Image of Glossy Black Objects. Surface Normal is given at a point chosen by the user.

navigation frame. The geometry of the surface in the local navigation frame is also known to be perfectly up and as such can be directly compared to the estimate.

Results for the DIRSIG simulations can be found in Section 5.2. In the next section, the linear polarimeter is tested against the simulations presented in this section.

4.3.2.4 Physical System Tests . The physical system test consisted of placing a glossy black painted plate on a turntable and rotating it through a series of β angles. Images were taken at nine angles between $10 - 80°$. The β angles used for this test are given to within $\pm 2°$. Although the room was well lit, a lamp was used as the primary source and was placed at a specular orientation from the camera. Average degree and angle of polarization values were determined for the specularly reflected section of the plate and plotted against the β angle. Figure 4.15 shows this plot along with the plots of the Fresnel and Shell models. This curve shows similar results to the DIRSIG simulation. It can be seen that the model fits well but that there appears to be a phase offset. This is likely due to misalignment of the linear polarizer in the rotation stage.
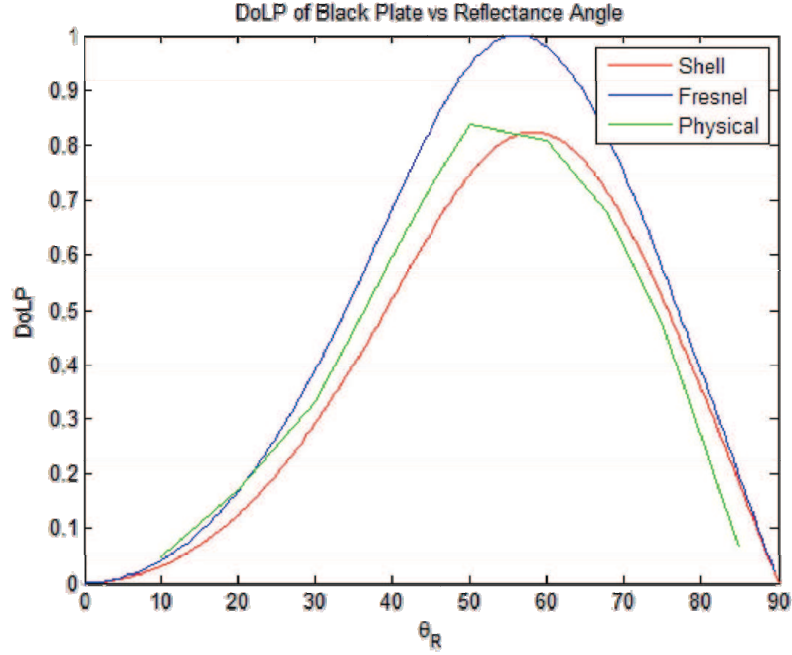
Figure 4.15:   DoLP vs β for the Fresnel Assumptions, Shell Target Model and Physical Measurements.

Figures 4.16 and 4.17 show typical angle and degree of polarization images for the physical system test. It is easy to see the importance of geometry in these images. The degree of polarization tends to fall off fast for non-specular angles.

Figure 4.18 shows a histogram of the DoLP along the vertical line in Figure 4.17. It can be seen that around pixel 300 the line is still on the black plate, but the degree of polarization starts to fall off as the specular geometry assumption is violated.

The same surface normal algorithm used in the DIRSIG simulation was tested with the physical system also. User defined points in the image and a β angle is estimated and surface normal computed in the frame of reference of the camera. Figure 4.19 shows an example of this surface normal on the glossy black plate.

A comparison of β angles was completed for each image in the simulation and physical system tests. Results of each of these tests are presented in the Section 5.2.

The surface normal of an object relative to the camera reference frame was determined using the method in Section 4.3.1 and was proven in the test described in
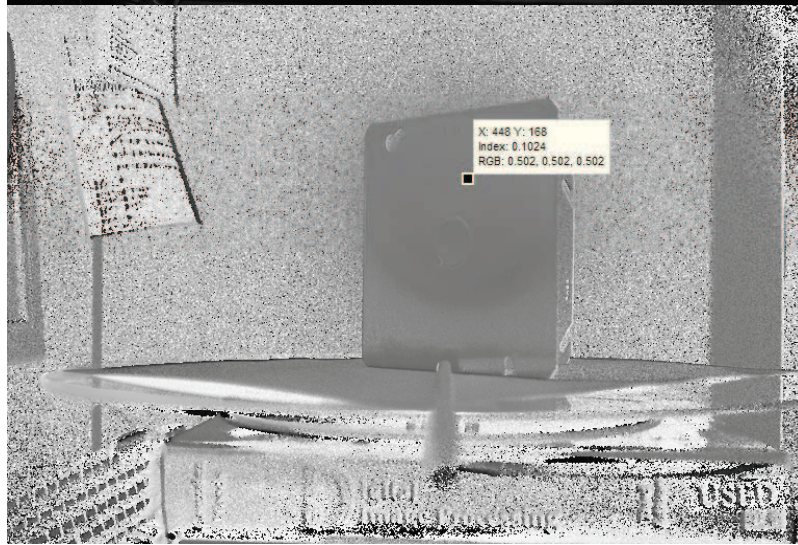
Figure 4.16: Example of an Angle of Polarization Image from the Physical System Setup.
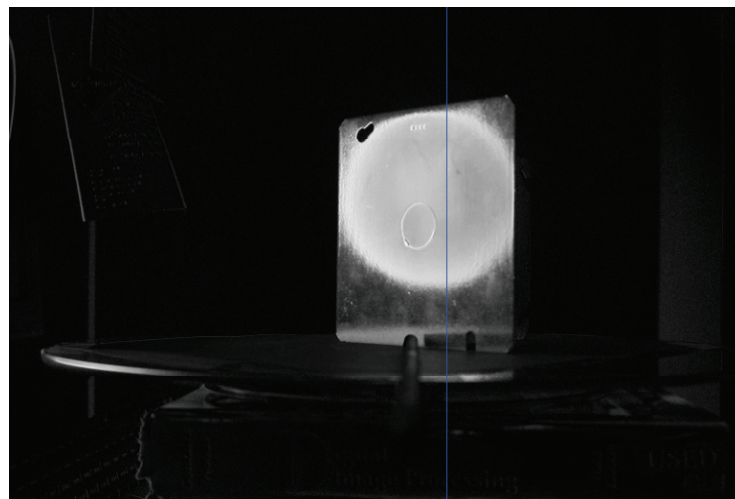
Figure 4.17: Example of DoLP Image from the Physical System Setup.
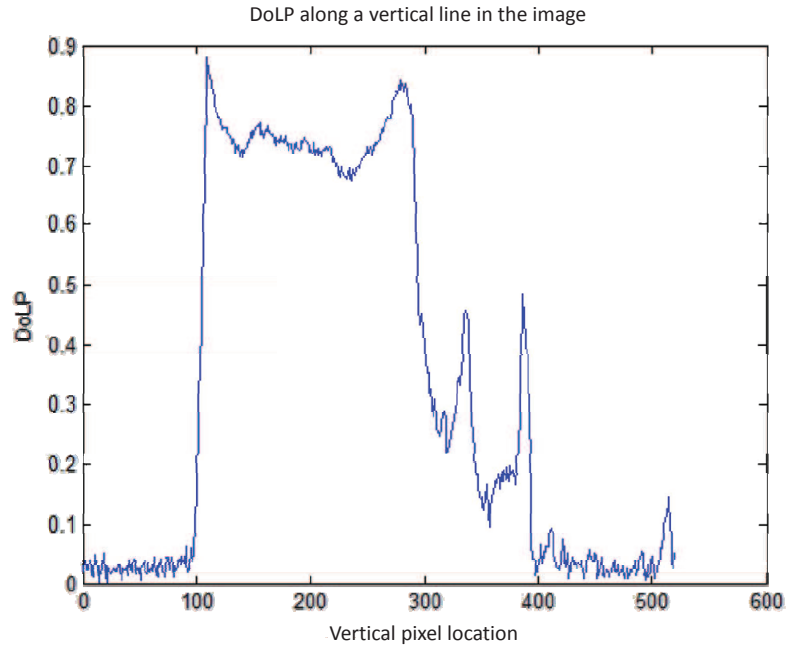
Figure 4.18:   Histogram of DoLP for the Line in Figure 4.17. The non specular falloff can be seen around the center of the image.
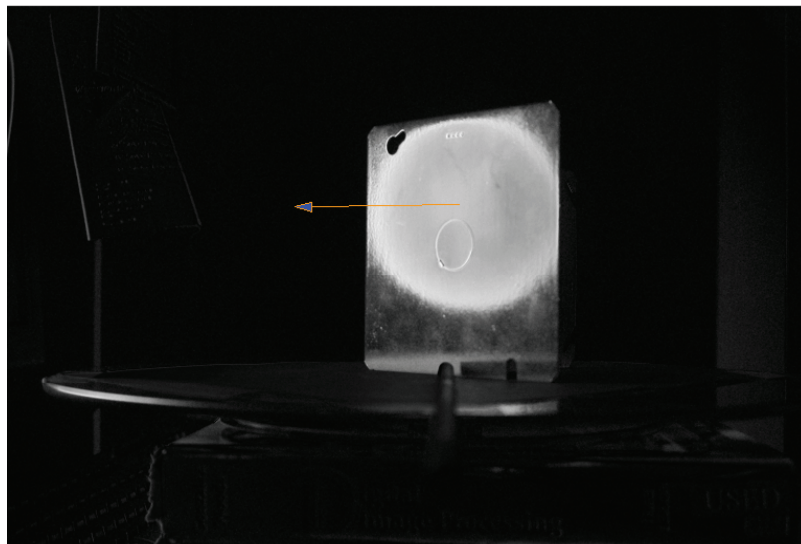
Figure 4.19:   Example of a Surface Normal Estimation Given a Point Chosen by the User.

Section 4.3.2. This is simply a relative orientation between the camera and surface. Therefore, if the orientation of the surface is known relative to a local navigation frame. The orientation of the camera can be found with respect to the navigation frame. Figure 4.20 shows an illustration of this relationship.
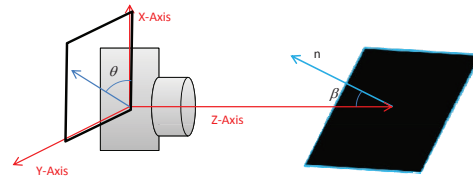
(a) Surface Orientation in Camera Frame
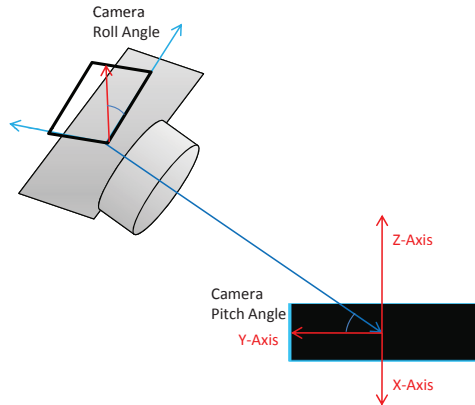
(b) Camera Orientation in Local Navigation Frame

Figure 4.20: Estimation of Surface and Camera Orientation. These images show how the estimation of pitch and tilt angles can be used to estimate the surface orientation in terms of the camera reference frame or the camera orientation with reference to a local navigation frame if the orientation of the surface is known relative to the navigation frame.

4.4 Estimation of Camera Orientation

In the last section, the receiver was assumed to contain the reference coordinate system. In this section, the reference coordinate system will be given in terms of the surface and the camera orientation will be estimated relative to that surface. Man-made surfaces tend to be constructed according to the Manhattan World Constraint [8]. By aligning the axes of the local navigation frame with the corridor of a hallway or room, many large flat surfaces tend to be found with surface normals pointing in cardinal directions. Using these assumptions about surfaces in the scene allows predictions to be made of the measurements of β and θ. Differences between these predictions and the measurements can then be used to update an initial estimate of the orientation of the camera.

The Unscented Kalman Filter approach found in Chapter II Section 2.7 is employed in this section to improve the estimate of the camera orientation. It determines which direction a surface faces and uses the measurements of β and θ to decrease the uncertainty in estimation of the DCM between the local navigation frame and the camera frame.

4.4.1 Algorithm .   The algorithm used in this section is a build up of algorithms used in all sections leading up to this. This algorithm requires a series of steps to be performed in order. This section describes the steps taken to perform a measurement update to estimate the camera orientation given only an image, taken in a well-lit environment, with a known material or smooth dielectric, having flat surfaces facing cardinal directions.

The UKF algorithm, described in Section 2.7, requires that the orientation of the camera be known with some uncertainty prior to incorporating the measurements. In order to get an initial estimate of the DCM for use in the UKF, the true measurements of the DCM are corrupted by random angle errors.

The flat surface grouping algorithm presented in Section 4.2 was used to find each side of the cube. The mean degree and angle of polarization of each side of

the cube were used to estimate candidates for β and θ angles. Recall that there are ambiguities in each of these angles and that there are two methods of determining β angle.

The candidates for surface normal orientation, β and θ are analyzed using the current estimate of the camera orientation. Predictions are made for what the surface normal vector in the camera frame would look like given these estimates and the estimates of the Shell parameters or estimates from the Fresnel model. A dot product is taken between the prediction and the measurements with the smallest dot product within a threshold revealing the most accurate surface normal orientation along with the correct β and θ angles.

Two models must be built in order to use the Unscented Kalman Filter approach, a measurement model, and a measurement uncertainty model. Sections 4.4.1.1 and 4.4.1.2 describe the models used for the measurement update portion of this UKF.

4.4.1.1 Measurement Model . The measurement model specifies the functional relationship between the true state vector and the observation. Given an a-priori state estimate, the measurement model can be used to predict the measurement realization. The predictions from this model will be compared to the best measurements of β and θ as shown in Chapter II Section 2.7.

In order to get a prediction of β and θ from the model, the orientation of the surface and the initial estimate of the camera orientation must be known.

Equations (4.5) - (4.7) were used for this prediction. They show that the surface normal in the navigation frame, $\boldsymbol{x}^n$, is rotated through the DCM of the estimate from the camera orientation, $\boldsymbol{C}_n^b$. This gives the surface normal vector in the camera frame, $\boldsymbol{a}^b$, in Cartesian coordinates. It was shown in Section 4.3 that the pitch and tilt angles of the vector of the surface normal are simpley the altitude and azimuth angles of the vector. A simple Cartesian-to-spherical coordinate transformation will yield the predictions for the β and θ angles.

$$a^b = C_n^b x^n \tag{4.5}$$

$$\beta = \arctan(\frac{\sqrt{a_1^2 + a_2^2}}{a_3}) \tag{4.6}$$

$$\theta = \arctan(\frac{a_1}{a_2}) \tag{4.7}$$

In these equations, the elements of the vector a are given as a_1, a_2 and a_3.

4.4.1.2 Uncertainty Model . The second piece that needs to be built is the measurement uncertainty model. In Section 4.3.2, the algorithm for determining the relative pitch and tilt angles between the surface and camera are shown. Uncertainties in these estimations are evaluated in Section 5.3.1.

Using these errors in estimation of β as starting estimates and tuning the Kalman Filter to give uncertainties aligned with true errors, the measurement error used for β was a Gaussian random variable with zero mean and a 6° standard deviation. These values represent the general uncertainty for a variety of materials, found using the MATLAB GUI described in Section 4.3.2.2.

The uncertainty in θ was determined through a transformation from uncertainty in angle of each intensity measurement to uncertainty in θ. This initial uncertainty was then tuned with the UKF to determine a zero mean uncertainty with 2° standard deviation.

Using these values for β and θ, the uncertainty matrix, R, used in this UKF was given as

$$R = \left(\begin{bmatrix} 6 & 0 \\ 0 & 2 \end{bmatrix} \frac{\pi}{180} \right)^2 \tag{4.8}$$

4.4.2 Test Setup .   Two tests were done to determine the effectiveness of the UKF algorithm. The DIRSIG simulation software was used to determine feasibility in a well-controlled environment. This test is discussed in detail in Section 4.4.2.1. Then, a test was done with the physical polarimeter in a real-world environment (Section 4.4.2.2).

4.4.2.1 DIRSIG Test .   In order to test the algorithm presented in Section 4.4.1, a single cube was given attributes of glossy black paint and placed in the center of the navigation frame with surface normals for each side of the cube pointing in the cardinal directions of the navigation frame. This cube was then surrounded with large flat surfaces with attributes of a totally reflective white paint in order to conform to the assumption that there is a well-lit environment with a source in each specular direction. A source was then placed in the positive z-direction, in an area that would reflect off of each of the white walls. The camera was initially placed in a position that would allow it to view three surfaces of the cube at even angles, and was then moved around to determine effectiveness of the algorithm at different orientations. Figure 4.21 shows the intensity, degree of polarization and angle of polarization images from the first camera view.

Recall from Chapter III Section 3.2 that DIRSIG outputs a set of Stokes images. These images were used to determine degree and angle of polarization using the method shown in Chapter II Section 2.4. Once degree and angle of polarization images were formed, the algorithm described in Section 4.4.1 can be used. Actual camera orientation is given in the DIRSIG simulations and can be compared with the estimations from the UKF.

4.4.2.2 Physical System Test .   A similar test was performed with the physical polarimeter. A glossy black cube was constructed using painted ceramic tiles. This cube was placed in the center of a set of bright white poster board sections. The source used was a desk lamp using a single, frosted 100-watt bulb and was placed close to the camera in a position that allows for as much lighting on the cube and

(a) Intensity   (b) Degree of Polarization   (c) Angle of Polarization
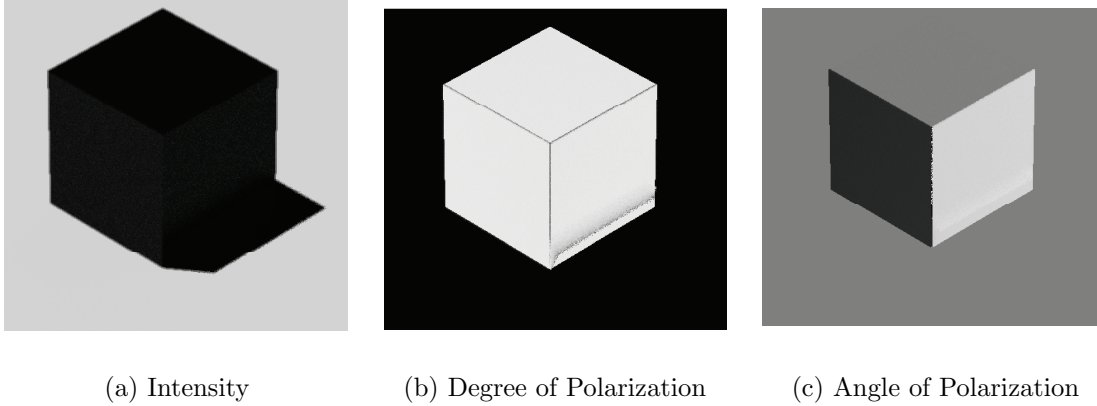
Figure 4.21:   DIRSIG Initial Orientation Images for Camera Orientation Estimation Test. These images show the intensity, degree and angle of polarization of the glossy black cube placed in a white box with an illumination source placed above. The initial orientation was set to give the camera the best view of three side of the cube in order to test the best case scenario for the UKF.

white walls as possible. Images were then taken from five vantage points similar to the DIRSIG vantage points. Figure 4.22 shows an image of the camera, source, cube, and white walls.

Degree and angle of polarization images were computed using the method described in Section 3.3. Figure 4.23 shows one example of the degree and angle of polarization images from the physical system.

In order to determine actual camera orientation for this test, a Vicon motion capture system was used [2]. The Vicon motion capture system works by using multiple cameras to track a fixed set of reflective targets. The targets are placed on an object and an object model is created. This model location and orientation can then be determined using similar methods to those described in Chapter II Section 2.3. Figure 4.24 shows a section of the Vicon motion capture area and the camera with reflective targets on it. To simplify the geometry of the truth comparison, the origin of the Vicon reference frame was set to be at the center of the cube, with cardinal directions pointing in the direction of the surface normals, just as in the DIRSIG simulation.

Figure 4.22: Camera orientation test setup in the Vicon lab. This image shows the camera, the light source, reflection panels and the glossy black cube test subject.

Each of the algorithms and tests done in this chapter was used to build on top of the previous tests in order to culminate in a set of models and measurements that were able to be used in a Kalman Filter in order to determine camera orientation. The constraints on the algorithm depend on the amount of information known about the scene. If a full set of Shell parameters are known, the only constraints are a well-lit environment with sources given in specular directions. If a full set of Shell parameters are not known, the constraints of a smooth dielectric material must also be upheld. Results for each of the above tests and an examination of errors that arise when these assumptions are violated are presented in the next chapter.

(a) Degree of Polarization

(b) Angle of Polarization
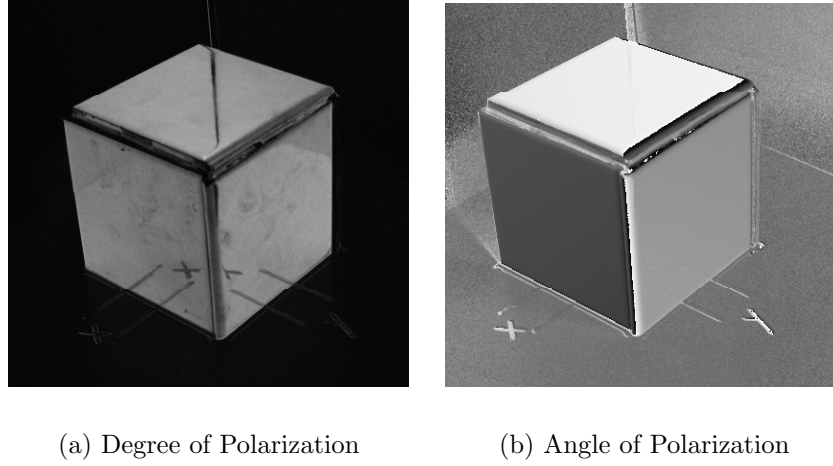
Figure 4.23:   Glossy Black Cube Images taken with the Physical System. These images show the degree of polarization and angle of polarization of the glossy black painted cube, imaged by the physical polarimeter, used in the camera orientation test. These images look similar to the DIRSIG images show in in Figure 4.21 .
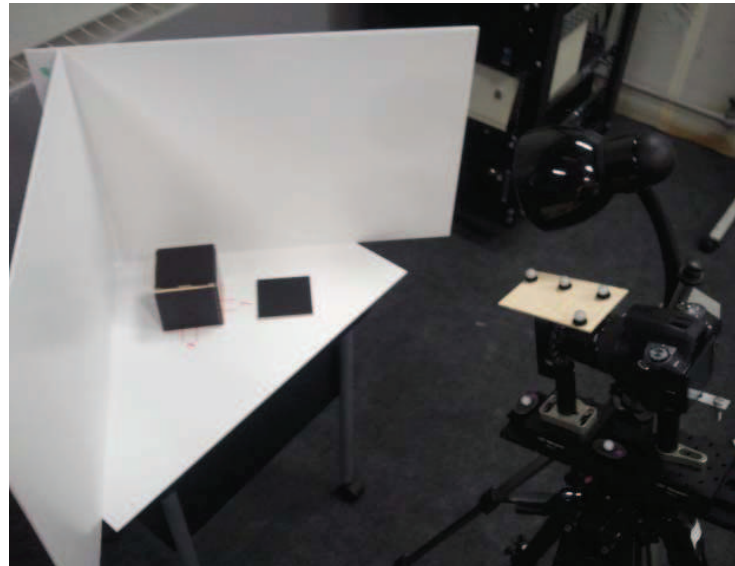
Figure 4.24:   Camera Orientation Setup In Vicon Lab. This image shows another view of the test setup. The reflective Vicon targets, placed on the camera, were used to determine actual camera orientation.

V. Test Results

There were several tests presented in Chapter IV, each with significant results presented in this chapter. Section 5.1 explores the limits of estimation of each surface parameter and relative geometry by showing the effect that each parameter has on the available measurements of intensity, degree of polarization, and angle of polarization. It also presents the results of attempts to estimate multiple parameters and explores the usefulness of each measurement by presenting the Jacobian and initial state change vector at a few key points. Section 5.2 then shows the results of estimating the surface geometry. Section 5.2 focuses on the limits of the Fresnel simplification algorithm presented in Section 4.3 but also shows an improvement in estimation of pitch angle if the full set of Shell parameters are known. Finally, Section 5.3 shows the results of the Unscented Kalman Filter implementation presented in Section 2.7.

5.1 Parameter Estimation Results

The set of Shell target model parameters is fit using a large set of measurements [31]. However, using these tools, it was determined that the estimation of a full set of parameters could only be done with a reasonable set of locations for navigation purposes if a multiple hypothesis testing method was used. Even then, this is only possible under specular geometry conditions. Given a full set of Shell target parameters, the surface geometry can be estimated well. Two sets of results are presented in this section. First, single parameters are analyzed individually to determine observability and envelopes of estimation. Then multiple parameter estimation results are presented in Section 5.1.2.

5.1.1 Single Parameter Estimation Envelopes . The algorithms and GUIs developed in Section 4.1 were first used to determine how well individual Shell parameters could be estimated. The results of this test show that the boundaries for estimation of a single parameter depend on the involvement of the parameter in the set of Shell target model measurement equations. This section describes the results of
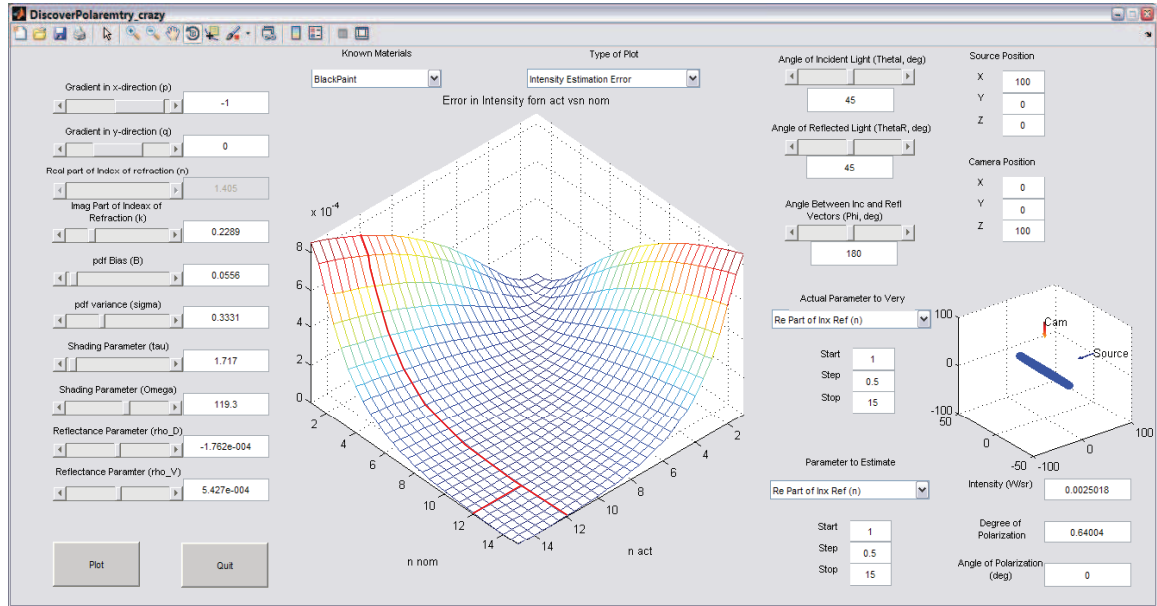
Figure 5.1: Error in Intensity as a Function of Actual and Estimated Surface Parameter. This illustration shows how the Levenburg-Marquardt algorithm will converge into valleys in the error curve. For the intensity measurement in this scenario, there are no incorrect valleys for the algorithm to converge into.

efforts to estimate single parameters with correct information about all other surface and geometry parameters. These results are key in determining which parameters can and can not be determined using a reasonable number of measurements.

The following figures show typical results for errors in the measurements as a function of a single parameter and are meant to demonstrate the usefulness of that measurement in the estimation of a single surface parameter. General limits of each parameter's starting estimates are presented in the following subsections.

Figure 5.1 shows that, using intensity as a measurement, there are no practical limits to the starting estimate of the real part of the index of refraction, n. For each individual slice of the actual parameter, there are no peaks or flat valleys in the graph with would cause errors in the estimate. However, this graph is misleading in the fact that the source is presented as a single point source at a precisely known location with a precisely known initial Stokes vector. If the source location or source Stokes vector is not known precisely, the estimates from the intensity alone are highly inaccurate.

96
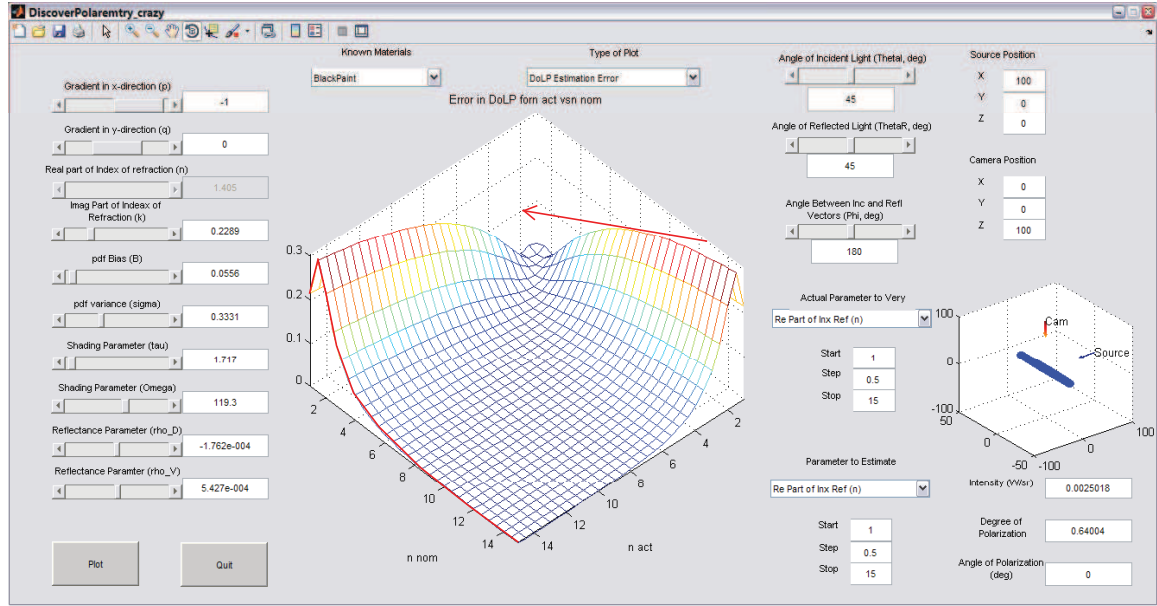
Figure 5.2: Error in Degree of Polarization Measurement as a Function of Actual and Estimated Surface Parameter. For this example, a starting estimate less than about 1.5 will cause the algorithm to converge into the wrong valley and give an incorrect estimate of the parameter. For some areas, starting estimates that are too far away will give initial change vectors that are too large. These large changes will jump over the correct alley and converge to an incorrect solution.

Figure 5.2 shows the error in degree of polarization measurements as a function of the same actual and estimated real part of the index of refraction. This graph shows that there is a peak for each actual value at around $n = 1.15$. This means that an initial estimate less than $n = 1.15$ will results in an incorrect estimate of the parameter. Also, for low values of the actual parameter, there are large errors in the measurement which can lead to an overshoot of the valley and will result in incorrect estimates as well. This was confirmed with the multiple parameter estimation GUI.

For intrinsic surface parameters, Figure 5.3 shows that using the angle of polarization measurements has no impact on the estimation. The lack of peaks and valleys means that the non-linear regression techniques will not move the initial estimate at all. This is intuitively obvious since the angle of polarization equations, shown in Chapter II, rely only on the geometry of the surface.
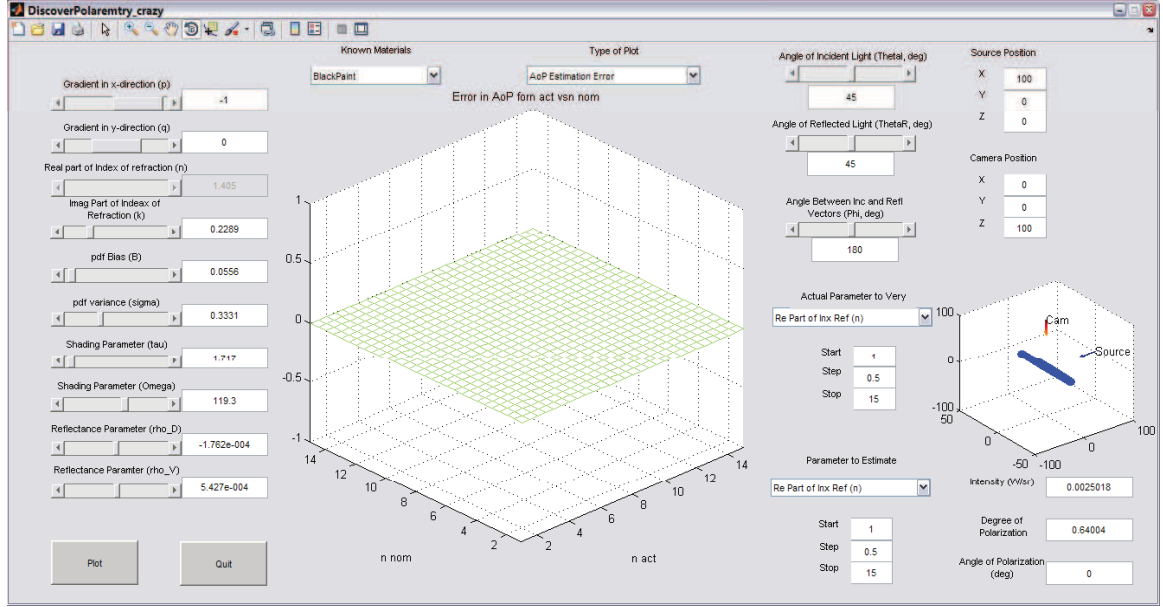
Figure 5.3:   Error in AoP Measurement as a Function of Actual and Estimated Surface Parameter. This figure shows that there is no effect on the angle of polarization as a function of this estimated and actual parameter. This is intuitive, since angle of polarization is simply a function of surface geometry.

Using this GUI and the multiple parameter estimation GUI, general limits for starting estimates of each individual parameter are given in the following subsections. Because of the possible uncertainty in source position or Stokes vector, the limits on starting estimates given in the subsequent sections are dependent only on the degree of polarization measurement.

5.1.1.1   Real and Complex Components of the Index of Refraction.

The components of the index of refraction, n and k, can be estimated individually with an accuracy proportional to the actual value. That is, for a dielectric material with $n \approx 1.5$ or $k \approx 0$, the starting estimate must be within 0.2 of the actual value in order to converge to the correct solution, and that convergence will generally yield an estimation within 0.01 of the actual solution. As these actual values grow (e.g., for a metallic surface), the requirement for a starting estimate is relaxed to a wider range. However, the overall accuracy of the final estimation, though reasonable, is not as accurate, generally presenting results within 0.1 of the actual value.

5.1.1.2 Surface Roughness Parameters. The components of the surface roughness parameters, B and σ, were more difficult to estimate. The B and σ components did not demonstrate observability using only the degree of polarization measurements. As a double check, it was shown that both parameters are very observable when the intensity measurement is taken into account.

5.1.1.3 Shading Parameters. The Shading parameters, Ω and τ, showed no observability at relevant geometries. Since these parameters only make contributions at very glancing angles, where the degree of polarization is already low, their contribution to the measurement is low, and therefore, they are impossible to estimate at relevant geometries.

5.1.1.4 Reflectance Parameters. Finally, the diffuse reflectance parameters, ρ_D and ρ_V, were generally the most difficult of the parameters to observe. These parameters greatly affect the overall diffuse return. It was determined that they were easier to estimate using brighter materials such as white paint, but that the starting estimates must be close to the actual parameter value. For darker paints, these parameters are very hard to estimate. However, common values for darker painted materials are on the order of 10^{-6} or smaller.

A summary of the results of each individual parameter test is shown in Table 5.1. It was determined that some parameters are not observable at relevant geometries, while others could be easily estimated. Correct estimation, however, requires correct knowledge of the other Shell parameters and the geometry of the system. In general, this is not a realistic scenario and either the full set of parameters are known or none of the them are known. The next section details the results of attempts to estimate a full set of Shell parameters and relative geometry. It also demonstrates how much must be known about a scenario in order to estimate unknown portions.

5.1.2 Multiple Parameter Estimation . It was determined using the multiple parameter estimation GUI from Section 4.1.2.2 that only specular geometries are

Table 5.1:   Table of Single Parameter Estimation Results. This table summarizes the results from the single parameter estimation tests.

Parameter	Notes
Complex Index of Refraction (n) and (k)	These can be estimated individually with a starting leniency and accuracy proportional to their actual value. They are easiest to estimate with only the degree of polarization measurement.
Surface Roughness Parameters (B) and (σ)	These parameters are more difficult to estimate with degree of polarization alone. However, for most situations there were very observable using both degree of polarization and intensity measurements.
Shading Parameters (τ) and (Ω)	These parameters were not observable under any relevant conditions or with any set of measurements. This is due to the fact that they have the largest affect at very glancing angles.
Reflectance Parameters (ρ_D) and (ρ_V)	These parameters are best estimated using the intensity measurement. There is some observability with the degree of polarization measurements, due to the fact that additional diffuse reflection will dampen the degree of polarization. However, using only the degree of polarization measurement requires a very close starting estimate.

useful in the estimation of surface parameters, only some of the parameters can be estimated using relevant geometries, and that fewer unknown states or more measurements generally led to closer estimates of the unknown states. This section describes the parameter sets that are best estimated. It also shows the results of the multiple hypothesis testing algorithm from Section 4.1.2.3 and its capability to capture a full set of Stokes parameters given a limited number of materials to chose from and an accurate geometry.

5.1.2.1 Interdependency of Parameters.   As a whole, the interdependence of free-fit parameters presented a challenge when trying to estimate multiple parameters simultaneously. The correlation between parameters and measurements generally caused the Jacobian space to present a minimization vector in an incorrect direction in order to more quickly reduce the residual of the measurement. This quickly led to an estimation which increased the error in a subset of the parameters. Whereas, in the last section, if all other parameters were maintained constant close to actual values, the estimate of the other parameter produced better overall results.

The multiple parameter estimation GUI was first used to estimate flat black painted material parameters and surface geometry with starting estimates consistent with a glossy black painted material and an accurate geometry. The sets of surface parameters for each of these materials can be found in Appendix A.

Figure 5.4 shows the multiple parameter estimation GUI results for trying to estimate flat black painted material parameters as well as surface geometry by starting with state estimates for a glossy black painted material, which has close parameters, and an accurate geometry. In order to understand how these errors come about, Table 5.2 presents the starting measurement influence matrix from the Levenburgh-Marquardt algorithm. This table shows how small changes in each of the parameters affect each of the measurements at the initial estimate values. The errors may also be understood by looking at Table 5.3, which shows the initial change in each state determined by the Levenburgh-Marquardt algorithm.

Table 5.2: Table of the Measurement Influence Matrix for intensity, degree of polarization and angle of polarization influence from Shell parameters near glossy black paint with unknown geometry. This table demonstrates how small changes in each parameter will affect each of the measurements.

Meas	x-grad	y-grad	(n)	(k)	(B)
Intensity	low	low	high	high	high
DoLP	none	none	low	med	none
AoP	none	none	none	none	none
Meas	(σ)	(τ)	(Ω)	(ρ_D)	(ρ_V)
Intensity	high	none	none	med	med
DoLP	none	none	none	low	med
AoP	none	none	none	none	none

Table 5.3: Table of relative initial state changes of the complete set of Shell parameters and geometry for glossy black paint. Notice the large changes in geometry even though those parameters are initially estimated correctly.

x-grad	y-grad	(n)	(k)	(B)
high	high	low	low	low

(σ)	(τ)	(Ω)	(ρ_D)	(ρ_V)
low	none	none	high	high

Since each parameter has a different range of values, the terms in these tables are given as qualitatively. Values are determined to be low if they are within 10% of the actual value or if they affect the measurement by less than 10%. Medium values range between 10% and 100%. High values are any values with an error greater than 100% of the actual value or change in measurement of more than 100%.

This table shows the percent change which would most quickly reduce the residual is achieved by changing the geometry estimates. It also shows that motion in the shading parameters will cause no effect in the residual, which is intuitive since it was shown in Section 5.1.1 that the shading parameters are completely unobservable in this geometry.

By reducing the state estimates to only observable states and providing a known geometry, Figure 5.5 shows some improvement in the estimates of the remaining sur-
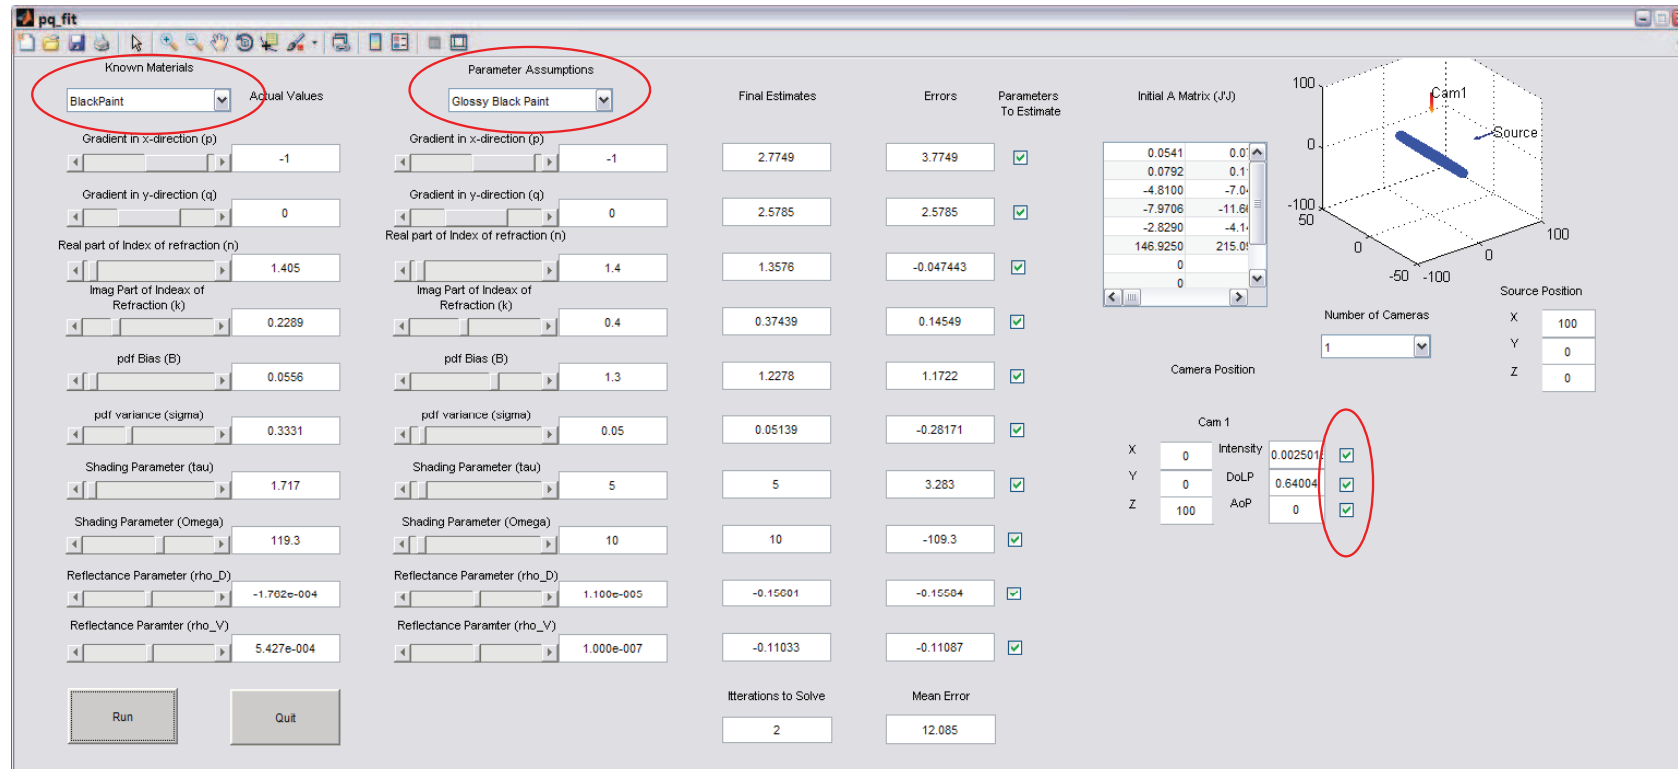
Figure 5.4: Multiple Parameter Estimation GUI Using an Estimate of Surface Geometry. The actual material is a flat black paint. It is being estimated as a glossy black paint with correct geometry, and all three measurements are being used at a single location. The result is major errors in all of the surface parameters and geometry.

Table 5.4:   Table of state change vector of the observable parameters of glossy black paint with known geometry. The initial changes in each state are much more reasonable compared to the initial changes for the unknown geometry example.

(n)	(k)	(B)	(σ)	(ρ_D)	(ρ_V)
low	low	low	low	high	high

Table 5.5:   Table of Jacobian for observable parameters of concrete with known geometry. This set of states shows a much smaller influence on the measurements for small changes in the parameters. This leads to large initial changes in parameter estimates in Table 5.6.

Meas	(n)	(k)	(B)	(σ)	(ρ_D)	(ρ_V)
Intensity	low	low	low	low	med	med
DoLP	low	low	med	low	high	high
AoP	none	none	none	none	none	none

face parameters. Using all three measurements yields an initial change in estimation shown in Table 5.4. These changes are much more reasonable but still exhibit large changes in estimates which are already close to the actual values.

However, if starting estimates are not close to the actual estimates, the set of unknown parameters can be thrown off in the wrong direction, as seen in Figure 5.6 and Tables 5.5 and 5.6. These figures show an example of estimating only observable parameters with a known geometry, but starting with a material that is not close to the actual material. In this case, the actual material is a flat black paint and the estimated material is concrete. Table 5.6 shows that initially the algorithm pulls the complex index of refraction the furthest, even though it is the closest to the actual value. This is because at this set of parameters, the intensity measurement has the largest impact on the state change and it shows that the index of refraction has little affect on the intensity measurement.

Because the changes to degree of polarization measurement as a function of each parameter are much closer than changes to the intensity measurement, only using the degree of polarization measurement will cause the initial change vector to be much smaller and not to exhibit such large jumps in close estimates.
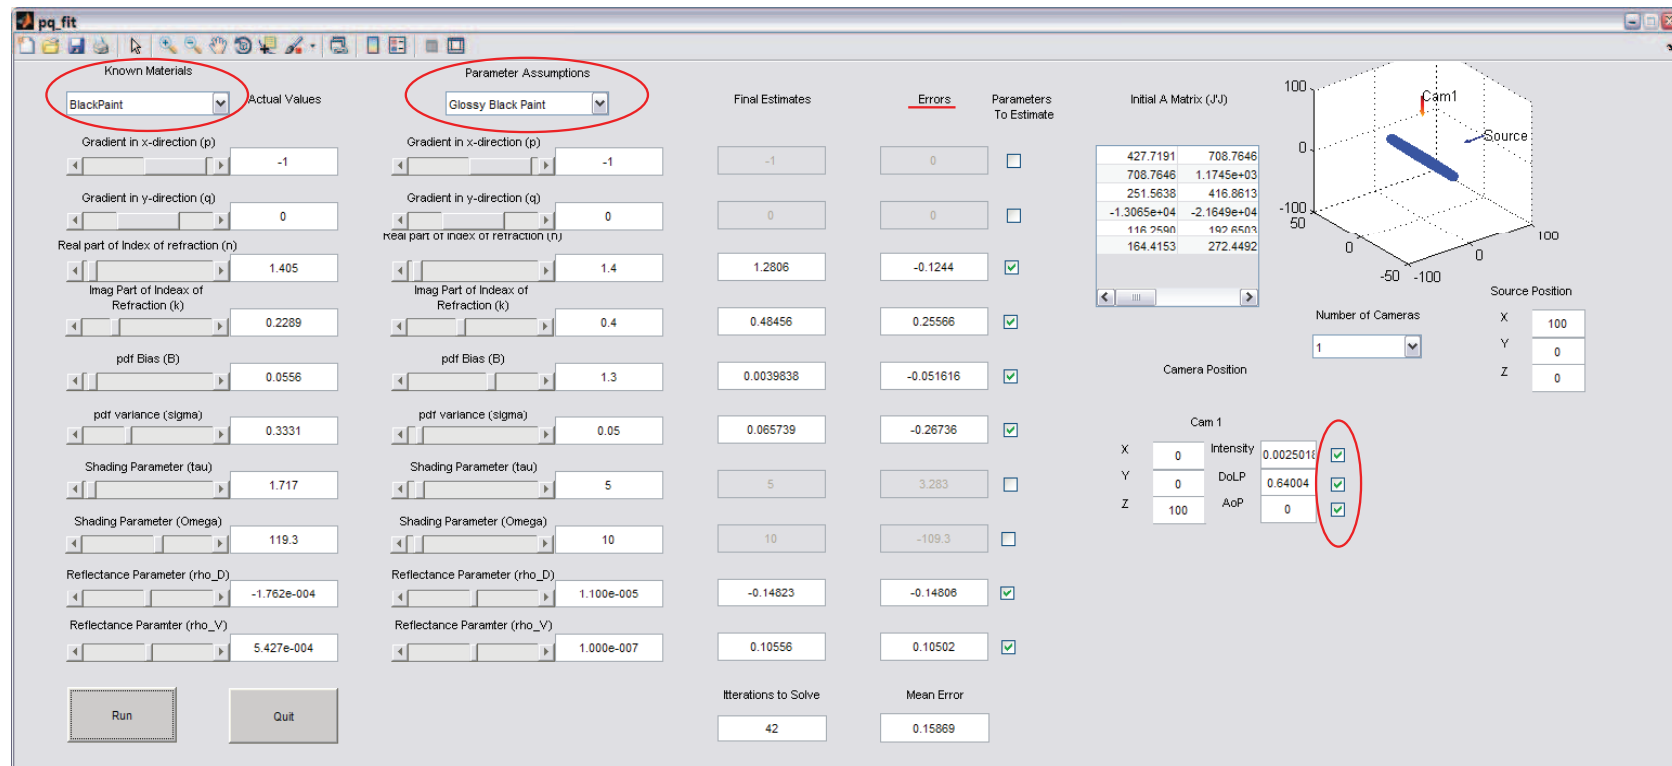
Figure 5.5: Multiple Parameter Estimation GUI Using Only Observable Materials and Known Geometry. Notice the dramatic reduction in error from the unknown geometry example.
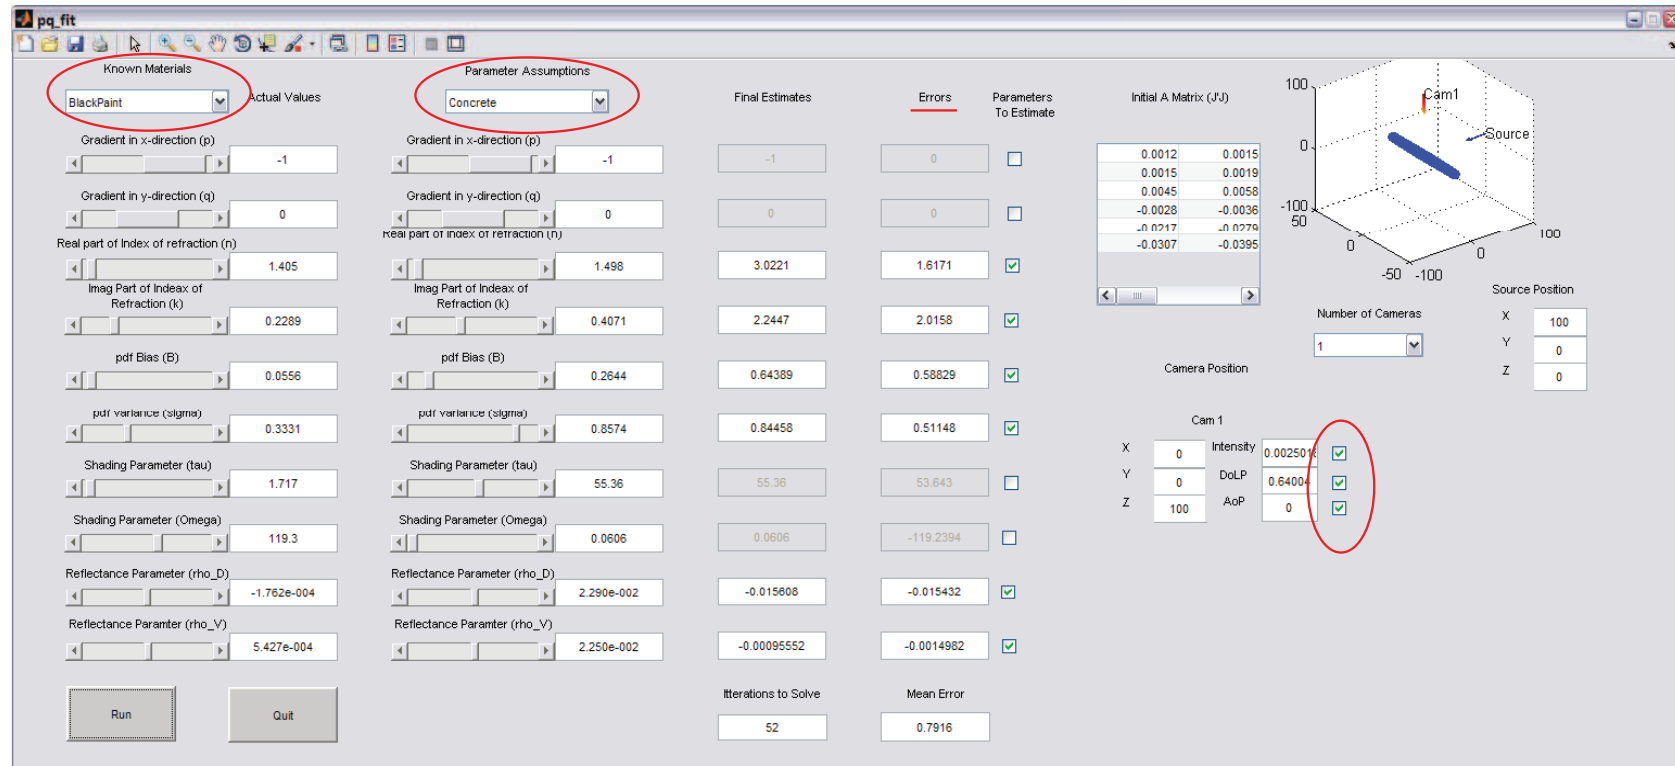
Figure 5.6:　　Multiple parameter estimation GUI using observable parameters and known geometry but with highly eronious initial estimates. The estimation of multiple parameters still requires reasonable starting estimates, similar to those determined in Section 5.1.1

Table 5.6: Table of state change vector of the observable parameters for concrete, which is far from the actual material, black paint. Large initial changes in the observable states cause overshoot of the correct solution and give erroneous final estimates for all parameters.

(n)	(k)	(B)	(σ)	(ρ_D)	(ρ_V)
high	high	med	med	med	med

As it was discussed in Section 5.1.1, the intensity measurement is highly suspect without absolute knowledge of the source. It is also shown in Figure 5.6 that the intensity measurement, even with absolute knowledge of the source, can have a negative impact on the estimate of observable parameters. The results of using the multi-parameter estimation GUI with only one degree of polarization measurement, in which the starting estimates are close and the user is only trying to estimate the observable parameters with a known geometry, shown in Figure 5.7, show a reduction in overall error of the estimates by about half. It is also intuitive, and can be seen by using this GUI, that multiple measurements of degree of polarization at different geometries will further reduce the errors in estimation of observable parameters so long as the geometry of the system is known.

It was shown in this section that not all of the Shell parameters can be estimated using a small number of measurements, that fewer unknown parameters and more measurements will lead to better estimates. These results mean that in order to estimate surface geometry using the Levenburg-Marquardt algorithm another method must be used to determine the full set of Shell parameters. The next section presents results from the multiple hypothesis testing GUI described in Section 4.1.2.3.

5.1.2.2 Multiple Hypothesis Testing . The results of the multiple hypothesis tests show that correct estimation of geometry plays a crucial role in estimation of surface type. When correct geometry is not known precisely, it was determined that multiple hypothesis testing leads to close, but incorrect, surface determination when only using the degree of polarization measurement. Using multiple measure-
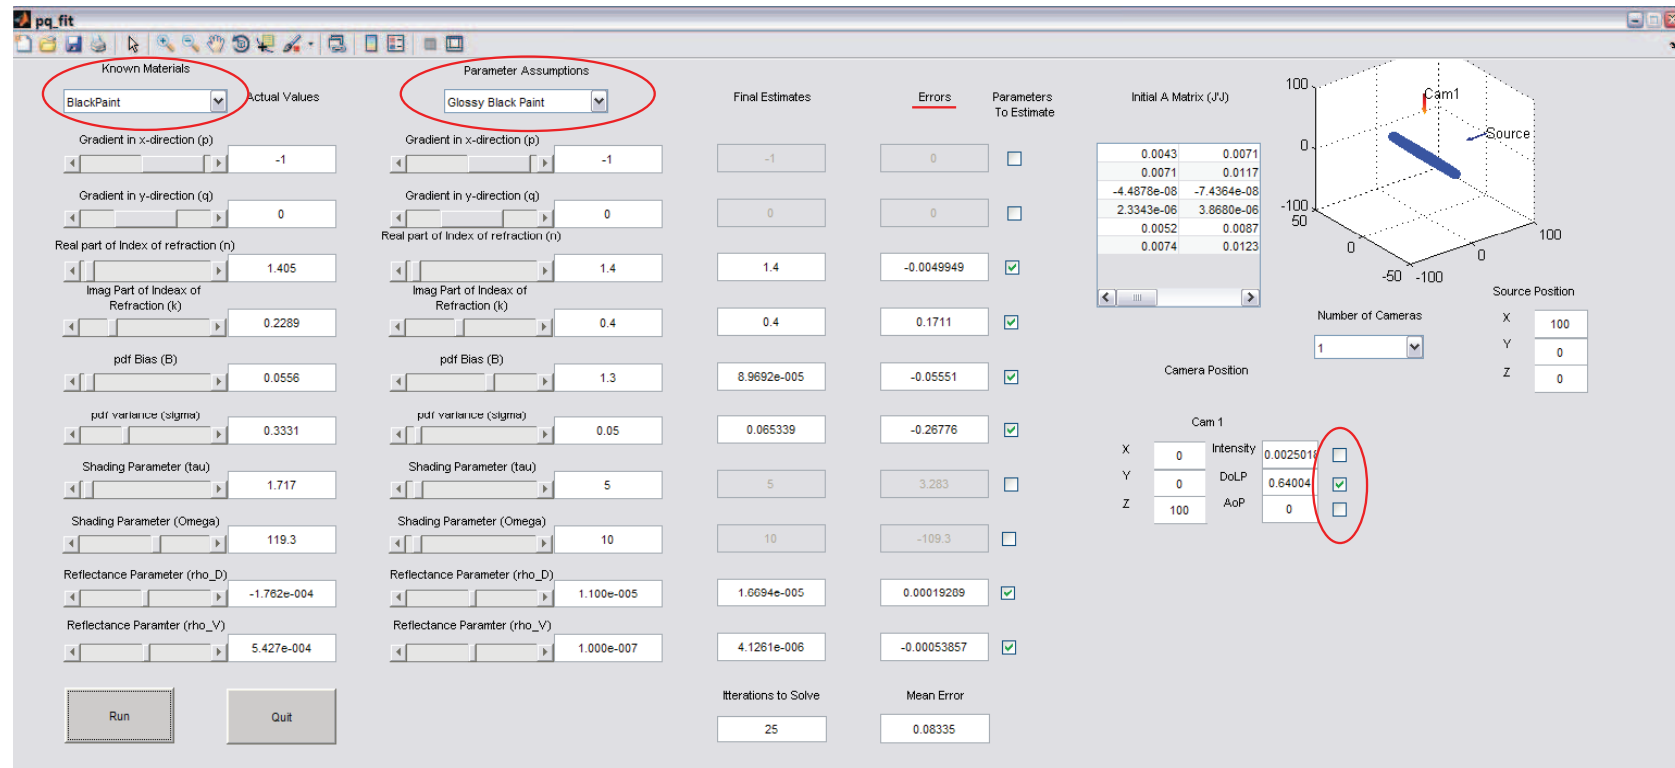
107

Figure 5.7: Multiple parameter estimation GUI using a known geometry, close initial estimate and only degree of polarization measurement. This example shows a reduction in final error of the observable parameters by removing the intensity measurement.

ments can mitigate this issue to an extent, but the final determination is that geometry must be known with a certainty relative to the specular spread of the material. It is also shown that incorrect initial estimates of surface parameters within $1 - 2\%$ of the actual value will still yield correct results for the surface estimate, given a correct geometry.

For targets without a specular spread, using only the degree of polarization measurement meant that the correct estimation of the surface is highly dependent on a correct geometry estimate. This can be seen in Figure 5.8 in which a glossy black paint is mistaken for a flat black paint with a small error in estimation of geometry.

Figure 5.9 shows that using the intensity measurement at off specular peaks causes errors in surface estimation given any error in geometry or surface parameters. This is due to the fact that specularities tend to cause glares and large changes in intensity at very particular angles. This can be seen by any observer that has noticed the glare off a window or a body of water.

Finally, Figure 5.10 shows the results of multiple hypothesis testing using a correct geometry but slightly incorrect parameters. This example shows a surface that is close to the glossy black paint material in the catalog, but with some differences in actual surface parameters. This GUI shows that multiple hypothesis testing will yield correct results even with errors in surface parameters of about 10%.

The results of the parameter estimation tests presented in this section show that it is only possible to estimate the full set of Shell parameters correctly, within the limits of a navigation scenario, using the multiple hypothesis testing method and that this method requires geometry to be well known. It also shows that the only usable geometry configurations are ones which present with specularities. These results were crucial in determining which constraints and assumptions could be made to simplify the determination of surface orientation. The following constraints are then practical and useful for determination of surface structure in an indoor environment. A surface
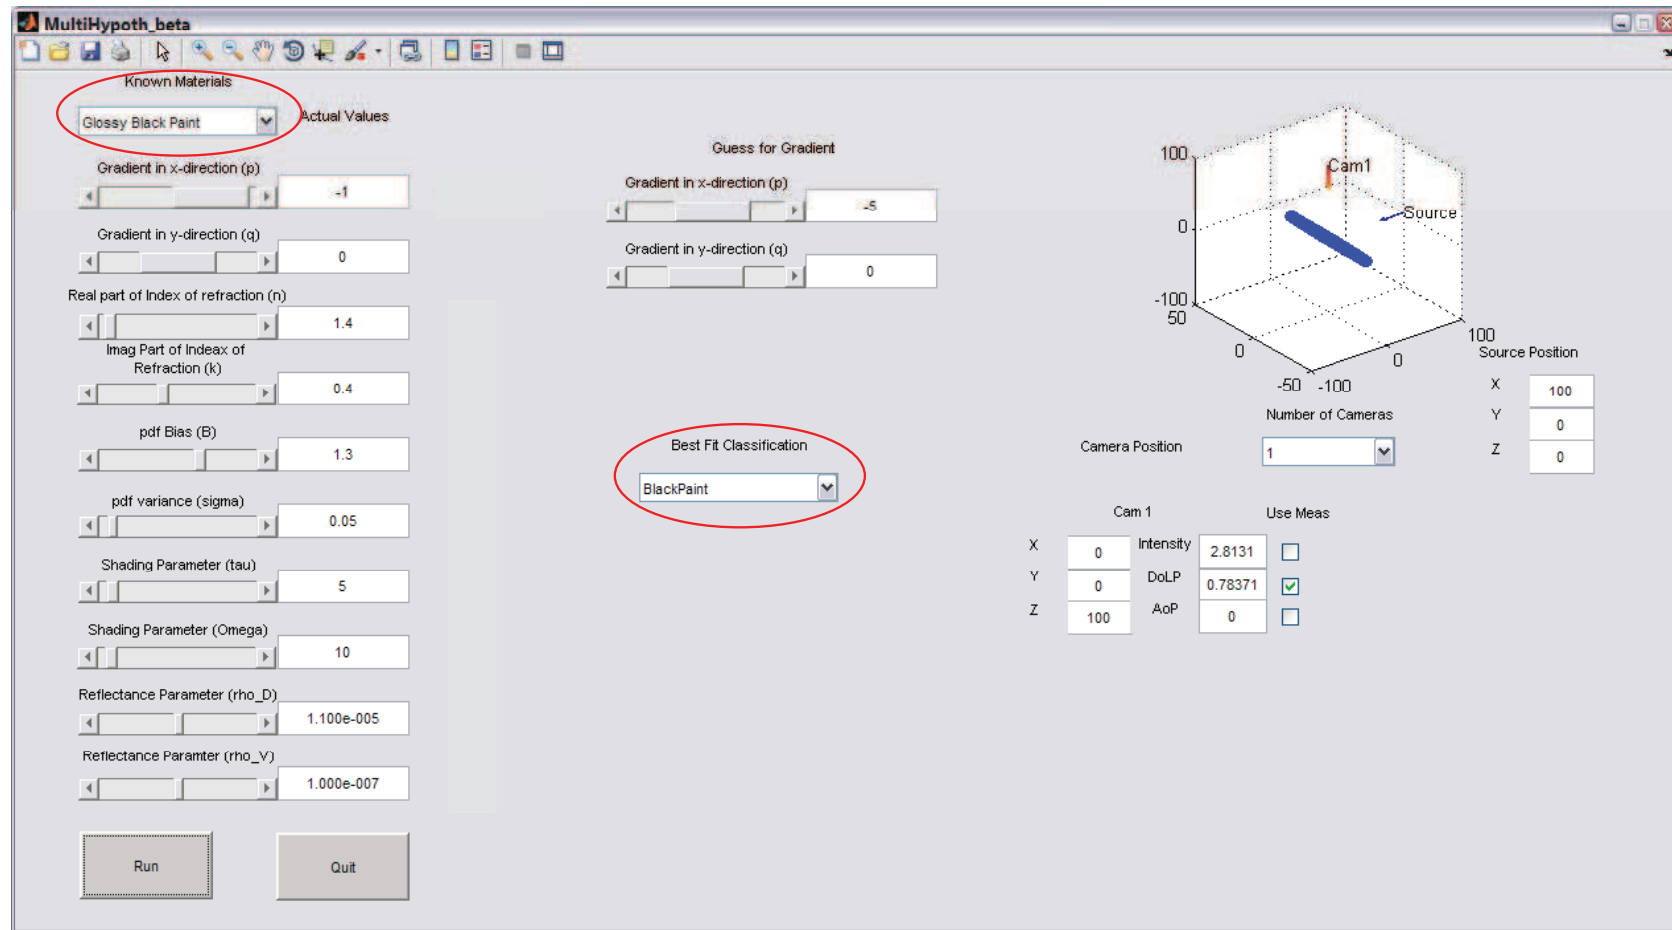
Figure 5.8:   Multiple Hypothesis GUI Showing the Results of Error in the Estimate of Geometry. The actual material, glossy black paint, is determined to be flat black paint.
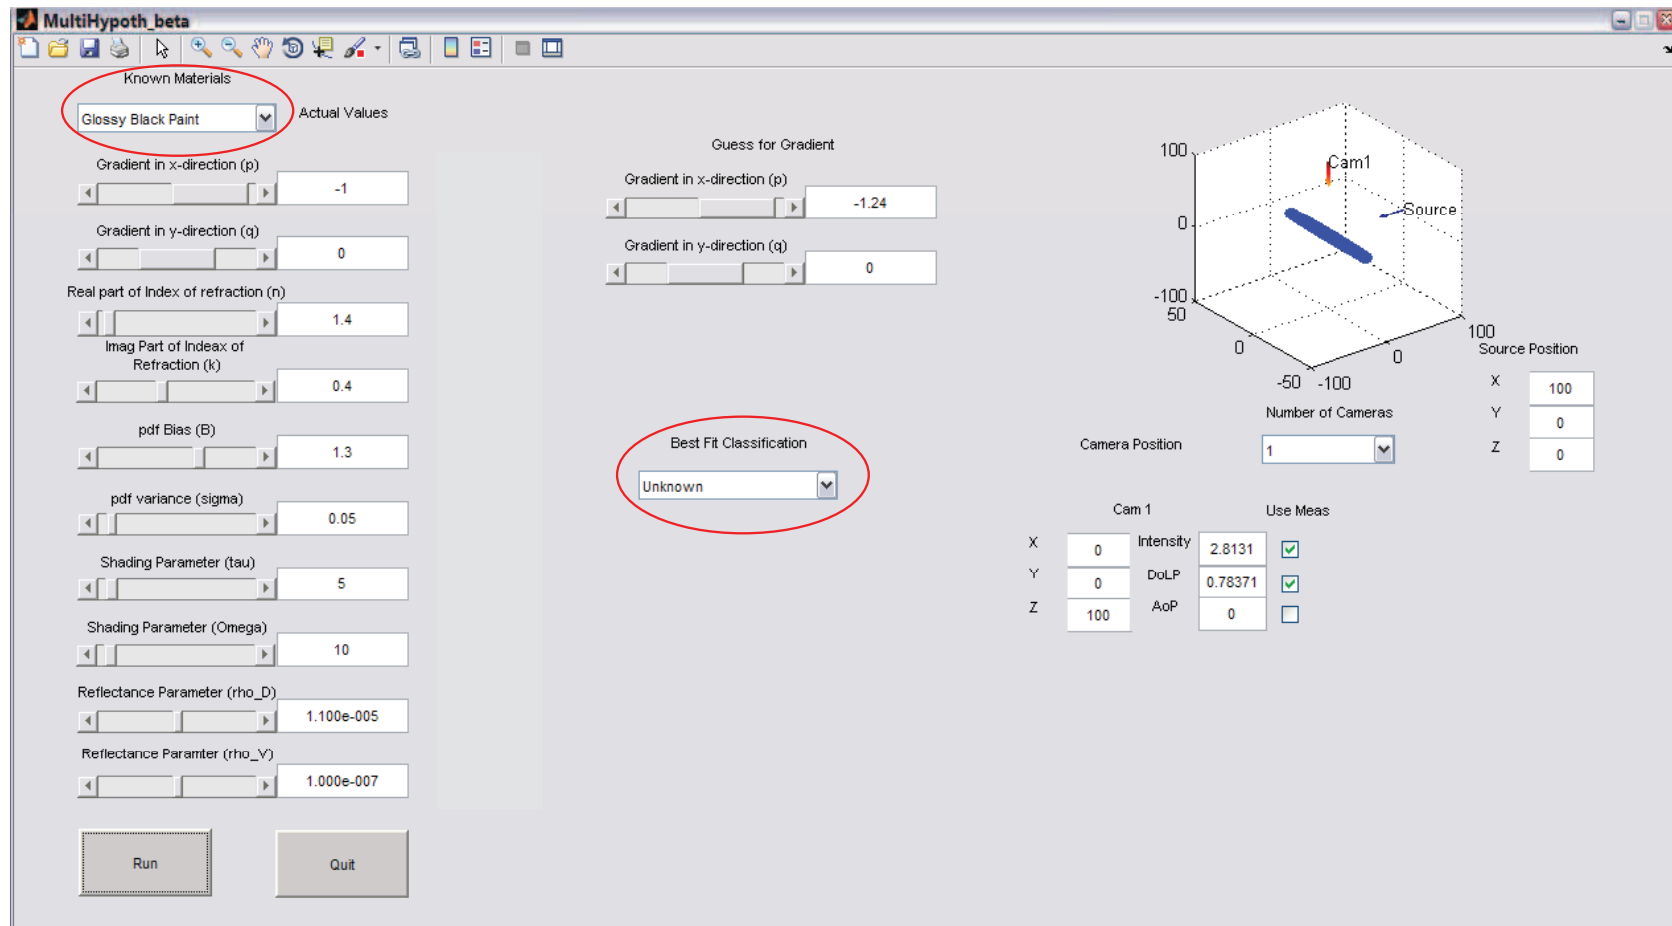
Figure 5.9: Multiple hypothesis testing GUI example showing how the use of the intensity measurement with errors in the estimation of geometry leads to an unknown material estimate.
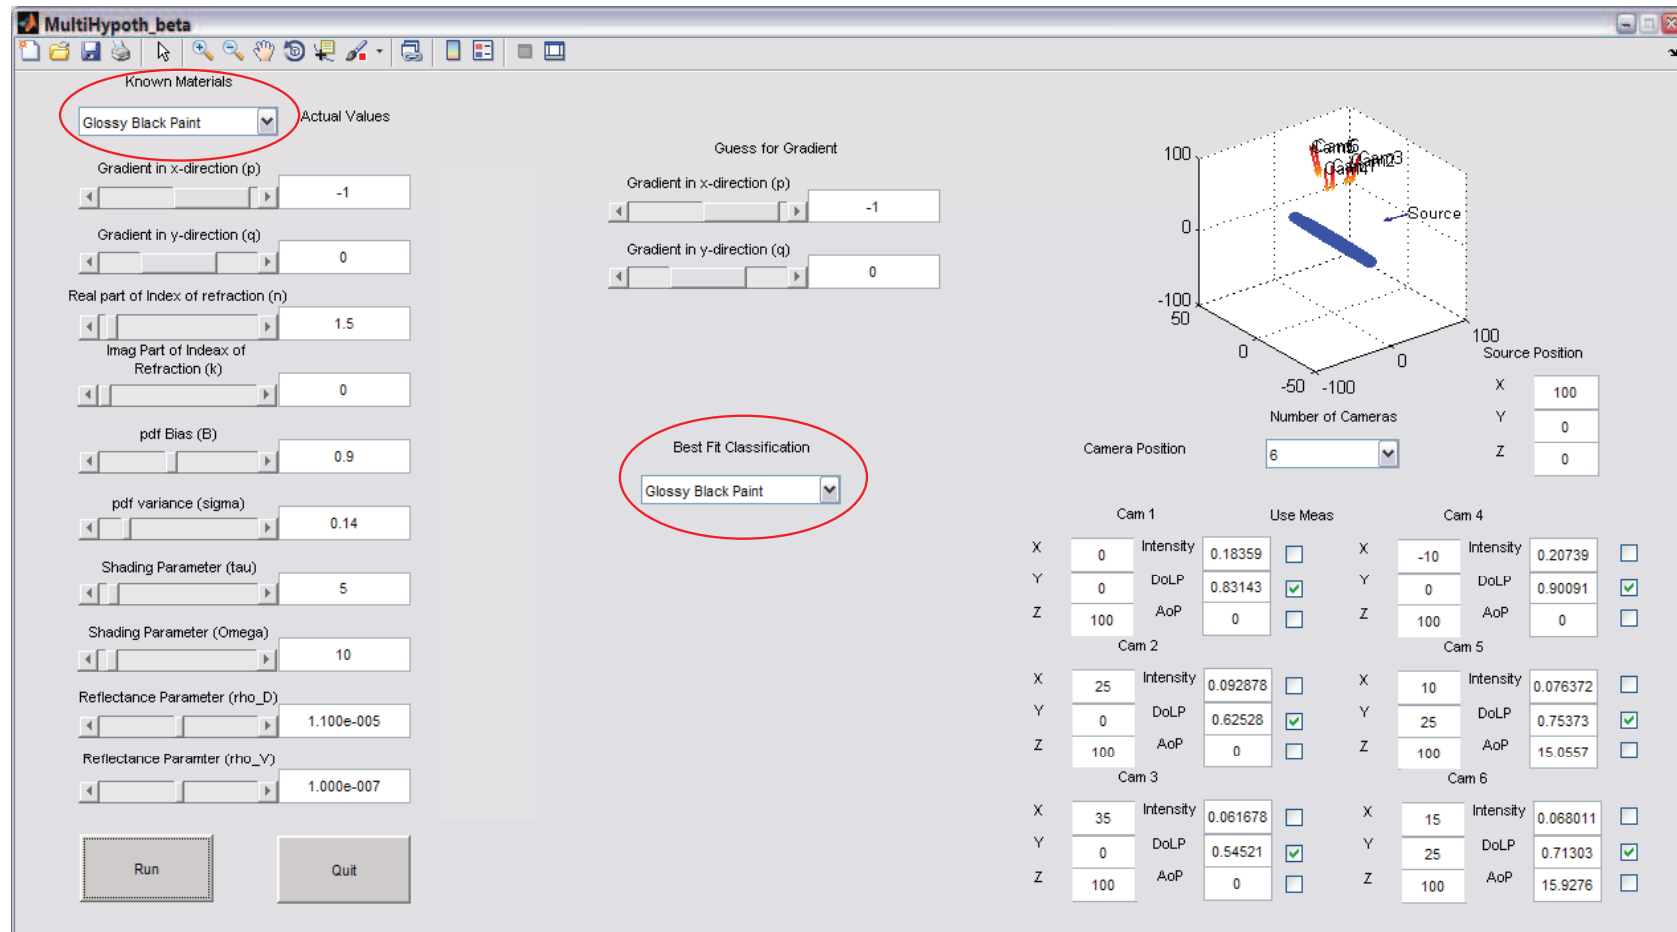
Figure 5.10:   Multiple hypothesis testing GUI showing the results of errors in the estimates of surface parameters, but with a correct estimate of geometry. The algorithm is fairly resilient with these types of errors, allowing parameter estimation errors of 10%.

must be be composed of a relatively smooth dielectric material and there must be a source which gives a specular reflection into the receiver.

5.2   Surface Orientation Estimation Results

Given the results from Section 5.1, it can be seen that, with a limited set of measurements, a full set of surface parameters can only be determined using a multiple hypothesis method and that a close estimate of the surface geometry must be known in order to use this method. Using these results, a set of constraints was implemented to alleviate the requirement for a full set of Shell parameters. These constraints are defined in detail in Section 4.3.

This section describes the results of comparing the simplified Fresnel algorithm to the Levenburg-Marquardt algorithm, which uses a full set of Shell parameters (Section 5.2.1). It then goes on to explore the limits of estimating the pitch angle given errors in the underlying assumptions (Section 5.2.2).

5.2.1   Results of Surface Orientation Tests .   Section 4.3 demonstrated the setup for a test which used the DIRSIG simulation software and the physical polarimeter to estimate tilt and pitch angles of a surface relative to the camera with the assumptions that the surface is composed of a fairly smooth dielectric material and that there is a source present in the specular direction. This section describes the results of that test. Recall that both the DIRSIG and physical system tests placed a flat plate of glossy black painted material in a specular geometry with the camera and then used the degree of polarization measurement to estimate the angle of reflectance of light off the surface, β.

The estimation of tilt and pitch angles in the DIRSIG simulation and the physical system were similar to the types of errors shown in the MATLAB simulation software. Figure 5.11 shows the actual and estimated β angles for both tests. The estimation of β is typically within 4°. For the DIRSIG simulation, it becomes worse near Brewster's angle, which is to be expected considering that the Fresnel estimation
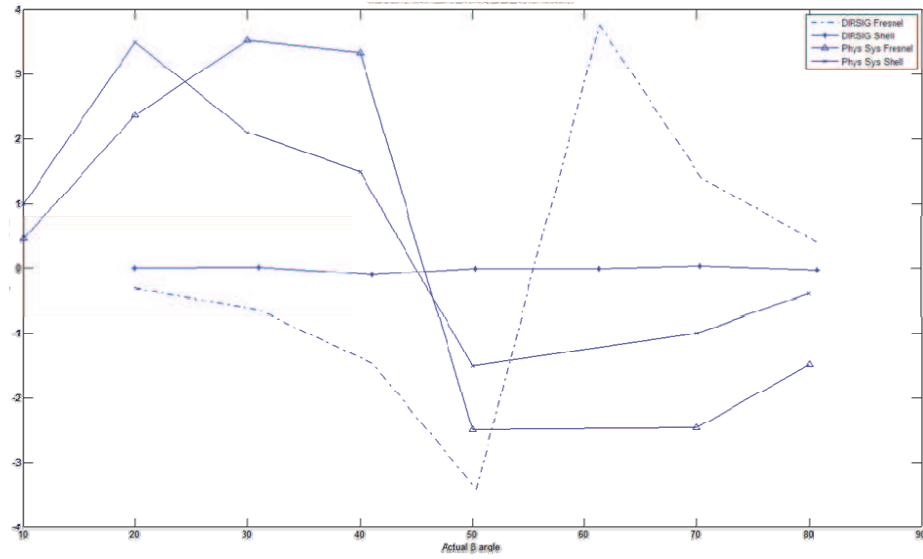
113

Figure 5.11: Results of the DIRSIG and Physical System Tests for Estimating β. This figure shows that the least amount of error occurs when the Shell target model is used with the DIRSIG simulation software. The Fresnel simplification model causes the highest errors around Brewster's angle for the DIRSIG simulation. The physical system shows more overall errors, but that there is little difference between the Shell model and the simplified Fresnel model, and that these errors are constrained to less than 4 degrees.

of the DoLP at that point tends to be much higher than the Shell target model values found around those angles. Errors for the physical system test were also under $4°$, but show a different shape than the DIRSIG simulations. This is likely due to the parameters of this particular paint sample being slightly different than the assumed values. However, unlike the DIRSIG simulation, the differences between the Shell target model estimation and the simplified Fresnel are much smaller.

This limited test gave rise to questions about how the estimation of β would fare for materials outside of the given assumptions. For example, could this method work with a metal instead of a dielectric, or just how smooth must the surface be for this method to work. The following section explores the limits of each individual parameter and the resulting error in the estimation of $/beta$ as a function of that parameter.

114

5.2.2 Envelope Analysis of Estimation Errors . Having verified the MAT-LAB simulation software, a set of tests was conducted to determine the extent of errors in estimation due to invalidations of the principle assumptions. Each parameter of the Shell target model, as well as the source-surface geometry was varied one by one and the error in estimate of the pitch was plotted against the actual pitch angle. For this example, the remaining parameters were left at the parameters for a glossy black paint. This example shows how a material that fits the assumptions may be varied by a single parameter and to what extent that parameter may be varied and still result in a reasonable estimation of the pitch angle.

The real part of the index of refraction was varied from 1 to 10 and results of the error in estimation are shown in Figure 5.12. These results show that the smallest errors in estimation occur around $n = 1.5$ and that errors increase as this value changes. This is due to the fact that the point $n = 1.5$ was used in the assumption in the simplification of the scenario. It can be seen from the graph that for a dielectric material, $1.4 < n < 1.6$, the errors tend to remain small, but that as n increases into ranges of, $n > 5$, the assumptions are severely violated and errors become unacceptable. These errors may be mitigated by using another assumption for the n and k values, given that those values for a particular material were known.

The complex component of the index of refraction, k, corresponds to the amount of light absorbed when an electromagnetic wave propagates through a material. A small k value then relates to a material with a low conductivity. This graph shows that for small values of k, the algorithm works well, but for values much more than 0.8 the errors start to spread into large geometries. These errors may also be mitigated for materials with a high k value by starting with an estimate higher than $k = 0$.

The σ parameter represents the spread of the specular probability density function (pdf) found in Equation (2.55). A lower $/sigma$ value represents a smoother, more specular surface. The graph in Figure 5.14 shows that as the spread of the specular pdf increases the error in β estimation increases. This is due to the DoLP

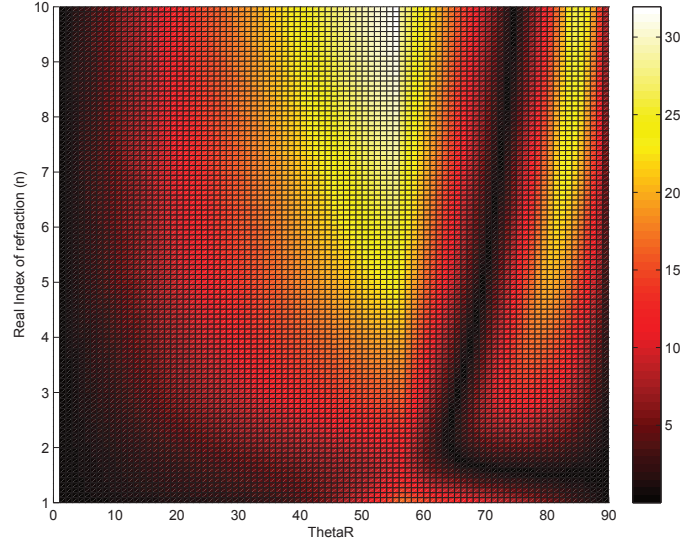Figure 5.12:   Error in Pitch Estimate vs Change in Real Component of Index of Refraction (n) and Actual Pitch Angle. As the index of refraction becomes larger than 1.6 the errors in the estimate of β become large.
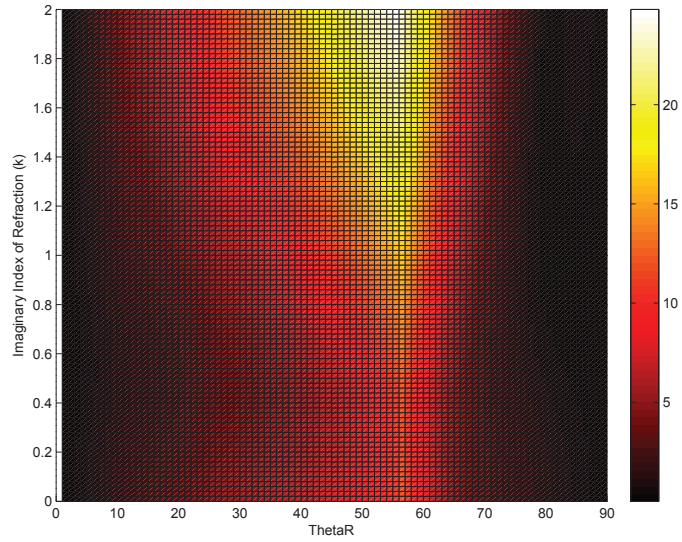

Figure 5.13:   Error in pitch estimate vs change in complex component of index of refraction (k) and actual pitch angle. Errors in the estimation of β become large as k becomes greater than 0.8.
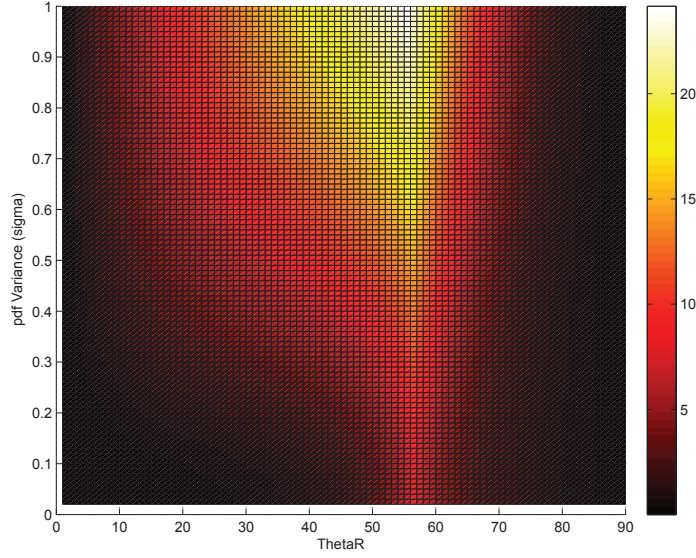
Figure 5.14: Error in Pitch Estimate vs Change in Specular pdf Spread (σ) and Actual Pitch Angle. Larger values of σ relate to a more rough surface. As the surface becomes more rough the assumptions break down and the errors in the β estimation become large.

being decreased in a particular direction because of this spread. This effect can not be mitigated by changing the estimates of the n and k values and shows that the assumption of a smooth specular surface should not be violated to any extreme.

The B parameter represents the bias in the spectral reflectance pdf and provides an overall magnitude adjustment to the pdf. Figure 5.15 shows that as B becomes larger, the error in β estimation decreases. This is due the the probability of spectral reflection increasing in direct correspondence. This allows for more of the overall DoLP to propagate past the interface and creates a DoLP curve close to the simplified assumption curve. As B is set closer to zero, the probability of spectral reflectance becomes almost zero and the estimate of the DoLP then becomes close to zero, which leads to large errors in the estimate of β.

The shading parameters, Ω and τ, are directly related in the shading equation given in Equation (2.57). This equation accounts for facets shadowed by other facets and can therefore diminish the reflected DoLP. Figures 5.16 and 5.17 show that these parameters have no effect on the estimation of β. They tend to correspond to depth
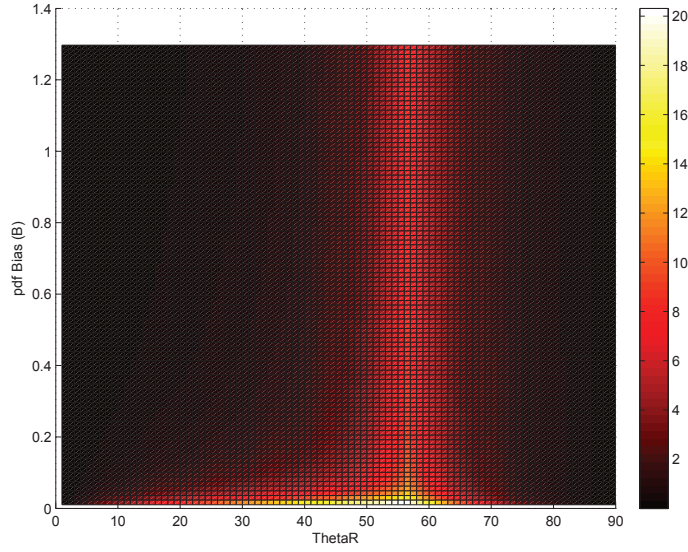
Figure 5.15: Error in Pitch Estimate vs Change in Specular pdf Bias (B) and Actual Pitch Angle. As the pdf bias decreases the degree of polarization measurement decreases and the errors in the estimate of β rise.

of surface roughness and only have effects at very glancing angles where the DoLP is already minimal.

The reflectance parameters used in the volumetric scattering equation given in Equation (2.58) can have a major effect on the estimation of β. The Fresnel reflectance equations assume only specular reflection in the computation of DoLP. Additionally, unpolarized light will add to the S_0 component of the Stokes vector and serve to reduce the overall DoLP. A material with a highly reflective diffuse component to the surface, such as white paper, would therefore not be ideal for this method. However, dark or glossy objects will still work well. The effects of the diffuse or Lambertian component, ρ_D, are shown in Figure 5.18. It can be seen that as this factor increases, estimates can quickly become unusable.

The effects of the volumetric parameter, ρ_V, are shown in Figure 5.19. This parameter shows a similar response to the diffuse scattering parameter. These errors can not be mitigated by simply changing the index of refraction estimates, however, they may be less of a problem in larger wavelengths where diffuse and volumetric reflectivity is not as apparent.

118
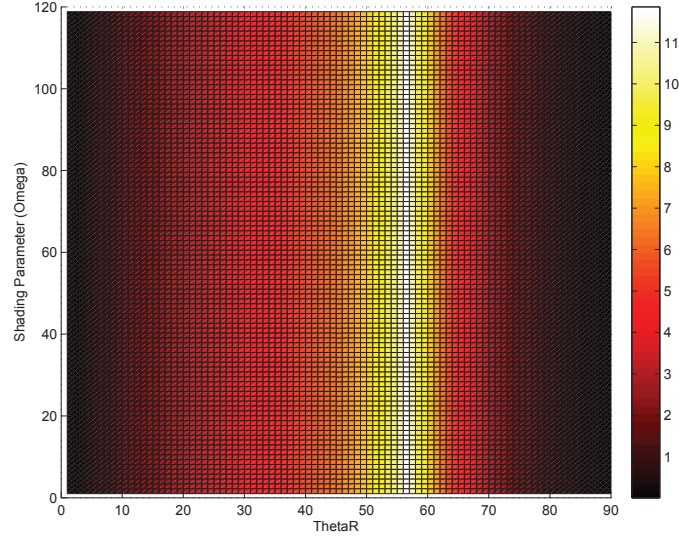
Figure 5.16:   Error in pitch estimate vs change in shading parameter (Ω) and actual pitch angle. For a given goemetry, changes in this parameter have no effect on the estimation of β.
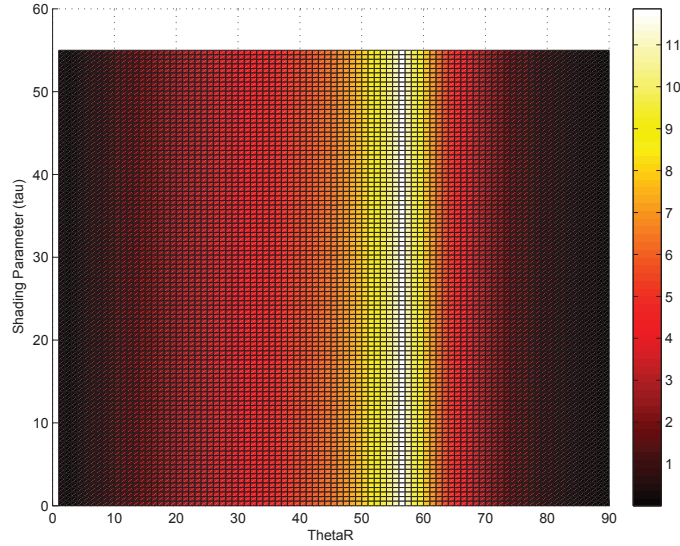
Figure 5.17:   Error in Pitch Estimate vs Change in Shading Parameter (τ) and Actual Pitch Angle. For a given geometry, changes in this parameter have no effect on the estimation of β.
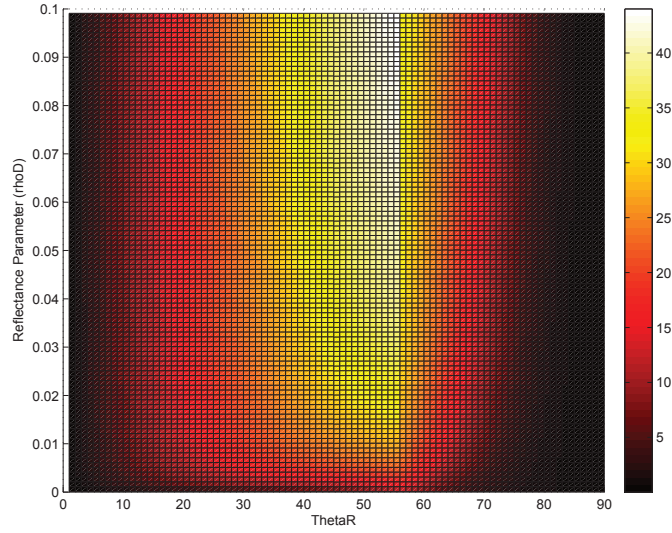
Figure 5.18: Error in Pitch Estimate vs Change in Diffuse Reflectance Coefficient (ρ_D) and Actual Pitch Angle. This shows that materials with a large diffuse reflection coefficient do not work well for this technique.
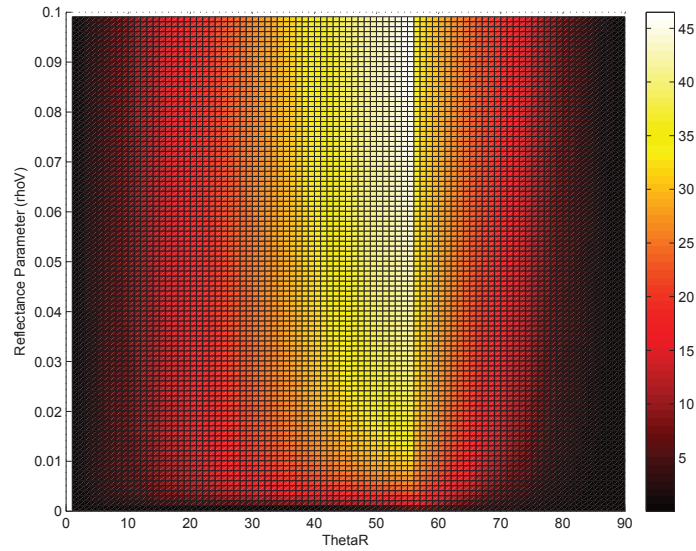
Figure 5.19: Error in Pitch Estimate vs Change in Volumetric Reflectance Coefficient (ρ_V) and Actual Pitch Angle. This shows that materials with a large volumetric scattering component do not work well for this algorithm.
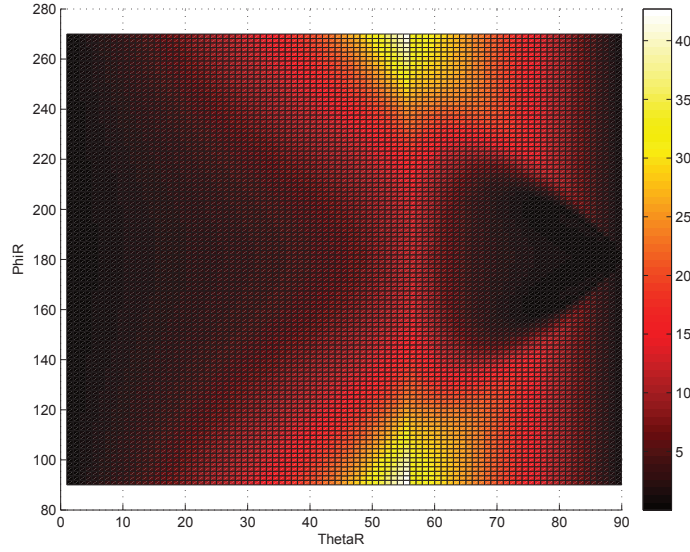
Figure 5.20:   Error in Pitch Estimate vs Change in Phase Angle Between the Source and Receiver (ϕ_R) and Actual Pitch Angle. This shows some leniency in the phase angle of the source for this glossy black paint example.

The specular assumption requires that a light source is directly in line with the receiver such that $\Delta\phi = 180°$. Figure 5.20 shows the errors presented when the light source is moved out of this phase. This graph shows that for this material, glossy black paint, there is some leniency in the line of sight parameter and that the algorithm can withstand errors in this geometry of $20-30°$ without much increase in estimation error.

The specular assumption also maintains that the incident light angle, θ_I, is the same as the reflected light angle, θ_R, and therefore the same as the estimation parameter, β. Figure 5.21 shows the effects of an error in the incidence angle. This graph shows that a difference between θ_I and θ_R simply changes the rising or falling slope of the DoLP curve and changes where the Shell target model curve intersects with the simplified Fresnel reflectance curve.

This section showed that, even without a full set of Shell parameters, the pitch and roll of a surface relative to the camera can still be estimated within a few degrees so long as the underlying assumptions of the Fresnel simplification model are not violated. It shows that the error in the pitch estimate was typically greatest around
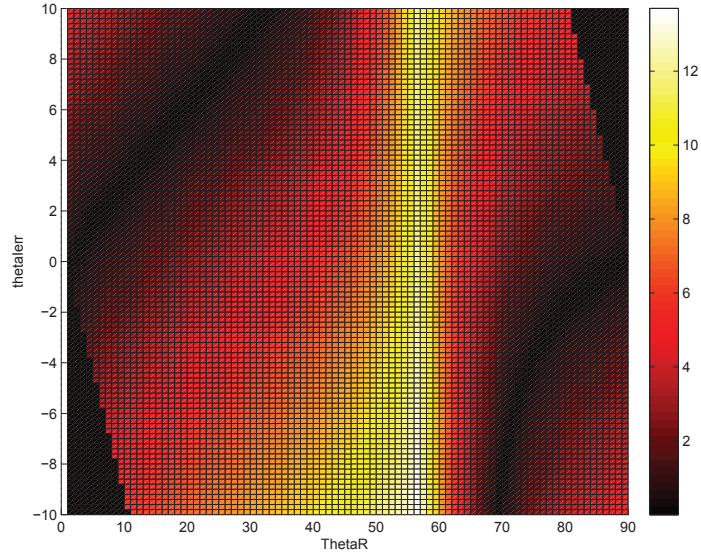
Figure 5.21: Error in Pitch Estimate vs Change in Difference Between θ_I and θ_R and Actual Pitch Angle. This shows that as long as there is a specular geometry, the error in the estimation of β is not affected by an error in pitch angle of $2 - 3°$.

Brewster's angle, due to the fact that most parameters of the Shell target model are used to dampen the degree of polarization measurement used in the estimation. However, it was shown that a full set of Shell parameters will result in a better estimation of the pitch angle. The next section describes the results of the Kalman Filter implementation and the estimation of the receiver orientation.

5.3 Camera Orientation Estimation Results

The results of the camera orientation test, though limited, show an improvement in certainty of camera orientation of roughly 25%. Improvement in estimation of camera orientation using degree of polarization measurements is a complex function of the number of measurements available, the material used in the estimation and the geometry of the scenario. The results presented in this section are based on an admitedly limited test set. However, given results presented in the rest of the chapter, it can be shown which scenarios should be more or less useful.

Although it was shown in Section 5.2 that it is possible to estimate the pitch angle of a surface without knowing the full set of Shell parameters, it was also shown

Table 5.7:   Summary of results for shape estimation tests using simplified Fresnel model.

Parameter Errors	Effects of Estimation of Pitch Angle
Difference between θ_i and θ_r	For the glossy black paint example, which has a moderate σ value, there is some lenience in the phase angle of the source. This is because the degree of polarization is spread along fairly broad area.
Offset in Phase angle ϕ	Within the specular geometry, the error in the estimation of β is not affected much by an error in incident light angle within $2 - 3°$.
Real part of index of refraction (n)	As the index of refraction becomes larger than 1.6 the errors in the estimate of β become large. This is due to the fact that the estimate of the real part of the index of refraction for a generic dielectric material is $n = 1.5$.
Complex part of index of refraction (k)	Errors in the estimation of β become large as k becomes greater than 0.8. This is due to the dielectric estimate of $k = 0$.
Surface roughness Parameter (B)	Smaller values tend to cause larger errors. As the pdf bias decreases, the degree of polarization measurement decreases and the errors in the estimate of β rise.
Surface roughness Parameter (σ)	Larger values of σ relate to a more rough surface. As the surface becomes more rough the assumptions break down and the errors in the β estimation become large.
Shading Parameters (τ) and (Ω)	For a given geometry, changes in this parameter have no effect on the estimation of β. There parameters were both shown to be unobservable in the last section.
Reflectance Parameters (ρ_D) and (ρ_V)	Materials with large diffuse or volumetric reflections dampen the degree of polarization and cause large errors in the estimation of β.

that errors in the estimation can become very high around Brewster's angle and as materials and geometry violate the assumptions given in Section 4.3.

It was also shown in Section 5.1.2.2 that if relative geometry is known within a few degrees, and a limited set of materials are found in the scene, that the full set of Shell parameters can be found using the multiple hypothesis algorithm. One of the benefits of using the Kalman filter approach in this section is that the camera orientation is always known, with limited uncertainty.

It was determined that using the Levenburg-Marquardt estimation algorithm for the β angle would be possible and would yield better results than the simplified Fresnel model. To this point, the β measurement uncertainty using the Levenburg-Marquardt algorithm has not been shown. Section 5.3.1 shows the results of the uncertainty in β estimation as a function of uncertainty in Shell parameters and geometry. The results of testing the filter with the DIRSIG simulations, described in Section 4.4.2.1, are shown in Section 5.3.2. Finally, the physical system test results are presented in Section 5.3.3

5.3.1   Uncertainty In Pitch Angle Estimates .   In Section 4.4.2, it was shown that the uncertainty in β measurement used in the Kalman Filter was 6°. This section shows how that number was determined and provides alternative solutions given a different scenario. The results for this section are also used to determine which scenarios might work better or worse in the estimation of camera orientation.

Recall from Section 4.4.2 that this test was conducted by using 300 particles with a 1% random error in Shell parameters at 1° intervals of β. The mean and standard deviation of the estimates of β were compared to the actual values and errors were plotted.

The results of five different materials are presented in this section. Figure 5.22 shows the error in β measurement for a flat black paint. Notice that the standard deviations of error are never more than 6°, but that there is a mean in the error around $\beta = 55°$. It has been shown before that most parameters in the Shell target model

124
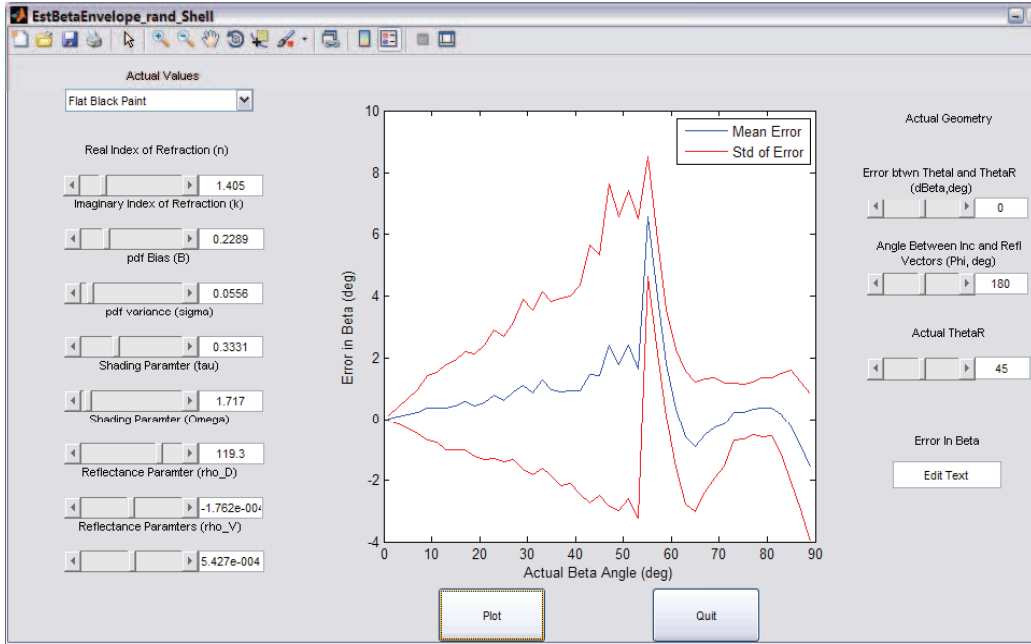
Figure 5.22: Mean and Standard Deviation of Error in Estimation of β for Flat Black Paint. This figure shows that the mean error of flat black paint increases near Brewster's angle, but that in general the standard deviation of error is less than 6°.

serve to dampen the degree of polarization measurement and that any geometry off of $\theta_i = \theta_r$ and $\phi = 180°$ will also result in a less than expected degree of polarization. This smaller than expected value causes the β measurement to fall to the left or right of Brewster's angle depending on the starting estimate. This figure shows that the flat black paint sample would be a good candidate for this algorithm if the mean error could be mitigated.

Figure 5.23 shows the results of a glossy black paint sample. This appears to be the best candidate for this algorithm. The errors in /beta are almost entirely less than 1° and the bias is close to zero except for a few degrees around Brewster's angle.

Figure 5.24 shows a tan paint example. This figure shows a much more stable material with a mean error close to zero and a standard deviation of error of 5° at most geometries. This material would prove to be very useful in the Kalman filter for β angles less than about 80°.
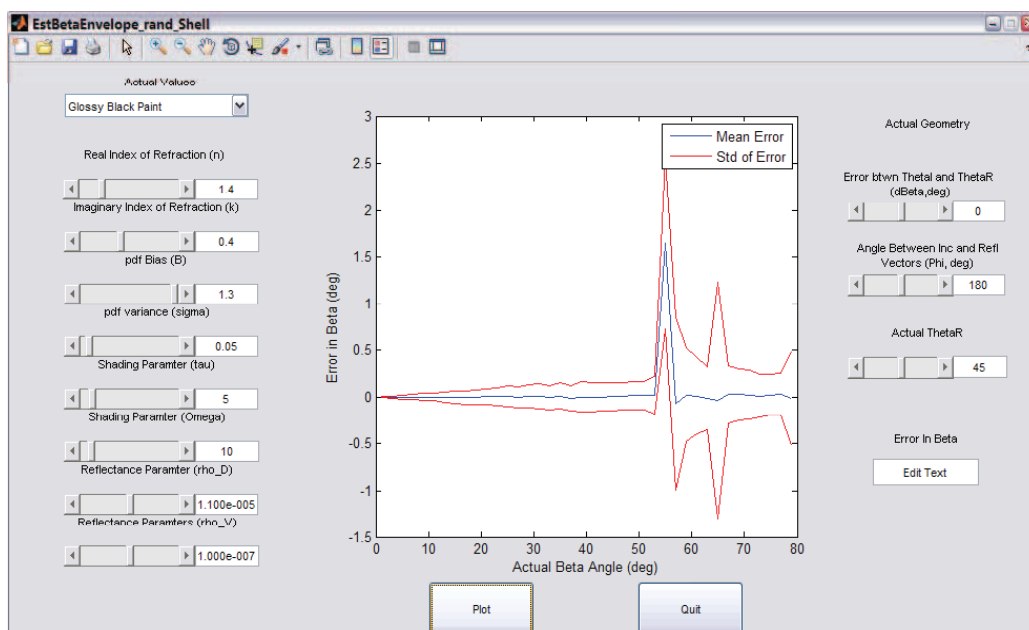
Figure 5.23:   Mean and Standard Deviation of Error in Estimation of β for Glossy Black Paint. This figure shows that glossy black paint would be a great candidate for the Kalman filter algorithm because the mean and standard devations of error in the estimation of β are typically very low, with the exception of right around Brewster's angle.
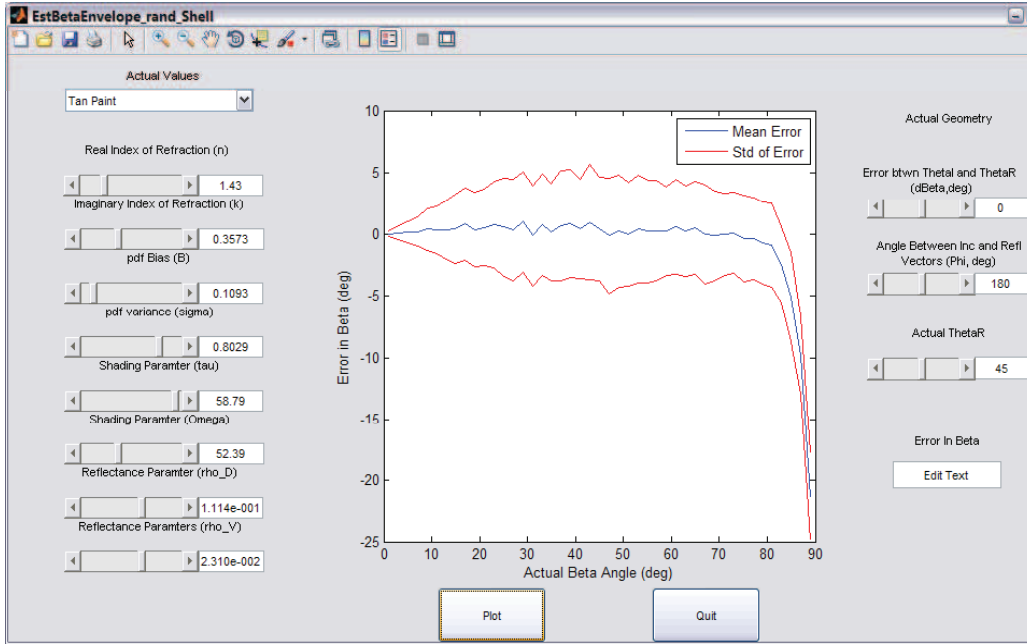
Figure 5.24:   Mean and Standard Eeviation of Error in Estimation of β for a Tan Paint Sample. This sample shows that it would also be a good candidate for the Kalman filter algorithm. It has a steady mean in error around zero and a standard deviation of error less than 5° for all actual β angles less than 80°.

A much worse example is shown in Figure 5.25. This is an example of a white paint sample. Recall from Section 2.5 that the white paint sample does not have a very broad specular peak and therefore any small error in geometry can lead to almost total loss of polarization. This would in turn cause the β measurement to fall far from the actual value, as seen in the graph. This figure shows that white paint would make a poor material for use in the algorithm.

Finally, Figure 5.26 shows the surprising results of a concrete sample. This image shows that this sample of concrete would make a good candidate for the algorithm. It shows almost no bias for most geometries and a standard deviation of error of less than 4°. The reason for this material acting so well is mostly due to the large polarization spread parameter ($\sigma = 0.85$). It will be shown in Section 5.3.2 that the error in geometry far outweighs the error in estimation of surface parameters, but because this material has such a large spread of degree of polarization it is less susceptible to such errors.
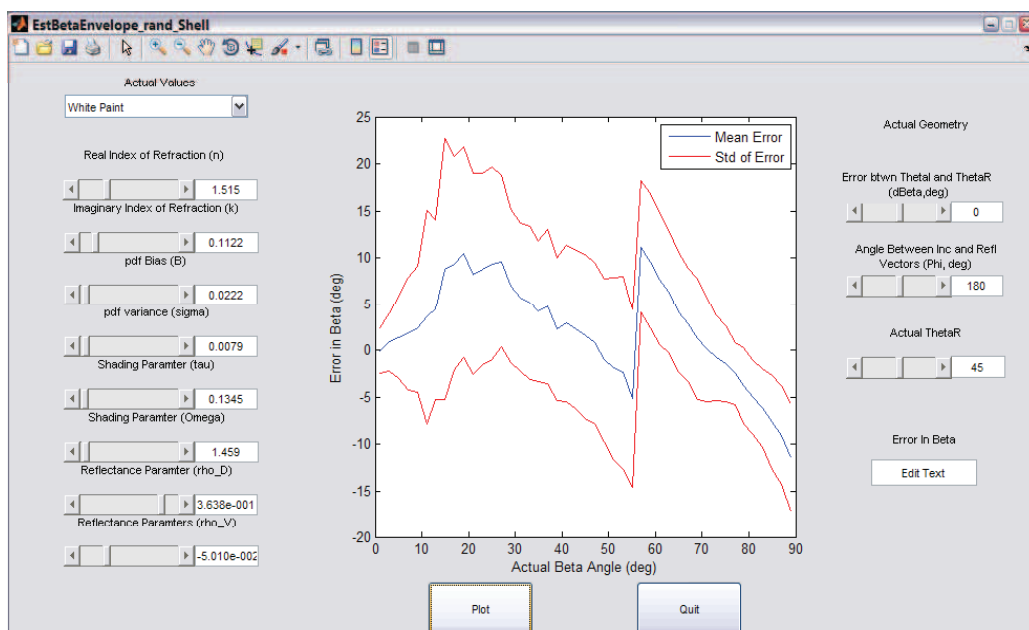
Figure 5.25:   Mean and Standard Deviation of Error in Estimation of β for a White Paint Sample. This figure shows that white paint would be a poor candidate for the Kalman filter algorithm. There are very large errors due to a very small spread in the degree of polarization off of the surface. A small error in angle thus has a higher impact on the degree of polarization measurement and causes a large error in the β measurement.
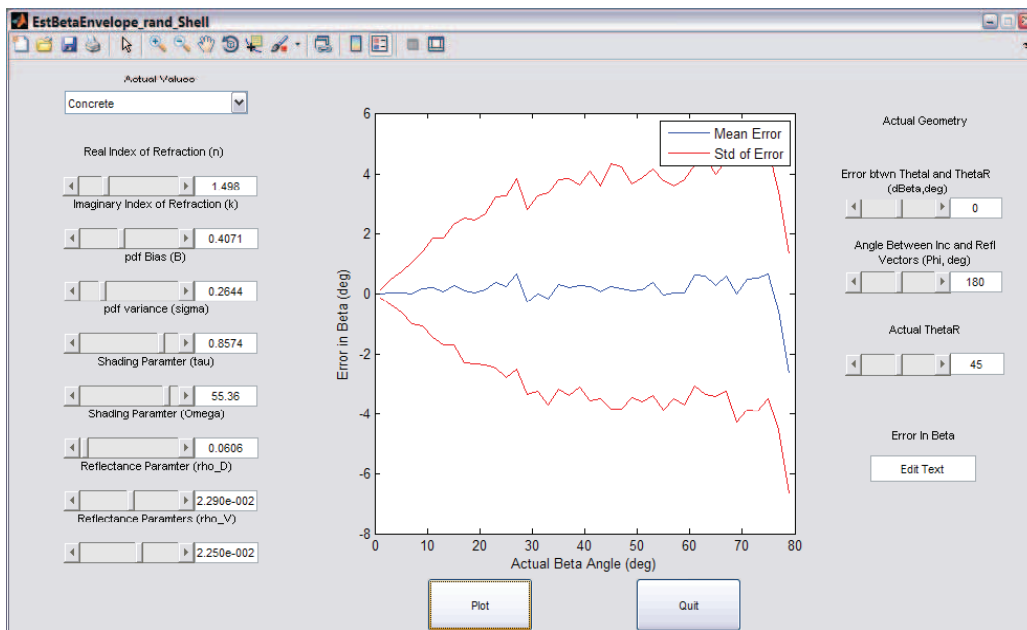
Figure 5.26:   Mean and Standard Deviation of Error in Estimation of β for a Concrete Sample. This figure shows, surprisingly, that concrete would make a good candidate material for the Kalman filter algorithm. It has a steady mean error around zero and a standard deviation of error less than 5°. This result is due to the fact that this particular sample has a large σ value which spreads the degree of polarization measurement. It will be shown in a later section that geometry plays an important role in the β measurement and this spread makes concrete more resilient to these types of errors.

These figures show that it may also be possible to predict the measurement bias and remove it. However, that test was left for future work. Given these results, a glossy black painted material was chosen for the DIRSIG and physical system tests. However, this β measurements uncertainty was left at 6° as a generalization.

5.3.2 DIRSIG Test Results .   The results of the DIRSIG simulations confirm the expected improvement from the Kalman filter model. Given a 6° uncertainty in the β measurement and a 2° uncertainty in the θ measurement, the expectation is a slight improvement in the x and y axes of the camera frame and a more pronounced improvement in the z axis of the camera frame.

Recall from Section 4.4.2.1 that the DIRSIG tests were set up at 25 altitude and azimuth angles between 15° and 75°. A glossy black painted cube was set at the center of the local coordinate frame with surface normals pointing in cardinal directions.

Each test sample was analyzed to determine which flat surfaces to use and measurements were limited to one per side. Though the algorithm was allowed to determine the orientation of each measurement in the local navigation frame, these decisions were compared to the actual orientations. Of the measurements used in this test, a very small number of incorrect orientations were chosen by the algorithm and these were at extreme geometries and extreme initial errors. Table 5.8 shows the altitude, azimuth and number of measurements used for each test.

Using these measurements, the first test was conducted to determine the measurement errors as a function of geometry alone. Figure 5.27 shows the results of this test. For each image, the correct orientation was initially fed into the filter. Errors in the measurements would then cause the correct orientation to be pulled in an incorrect manner. This figure shows that error caused by geometry alone is generally less than the filter expects. However, this is not the only source of measurement error to be tested.

Table 5.8: DIRSIG test conditions. This graph shows the randomized altitude and azimuth conditions for the DIRSIG tests. For some geometries, measurements could not be taken on some sides of the cube. The number of measurements relates to the final uncertainty in camera orientation.

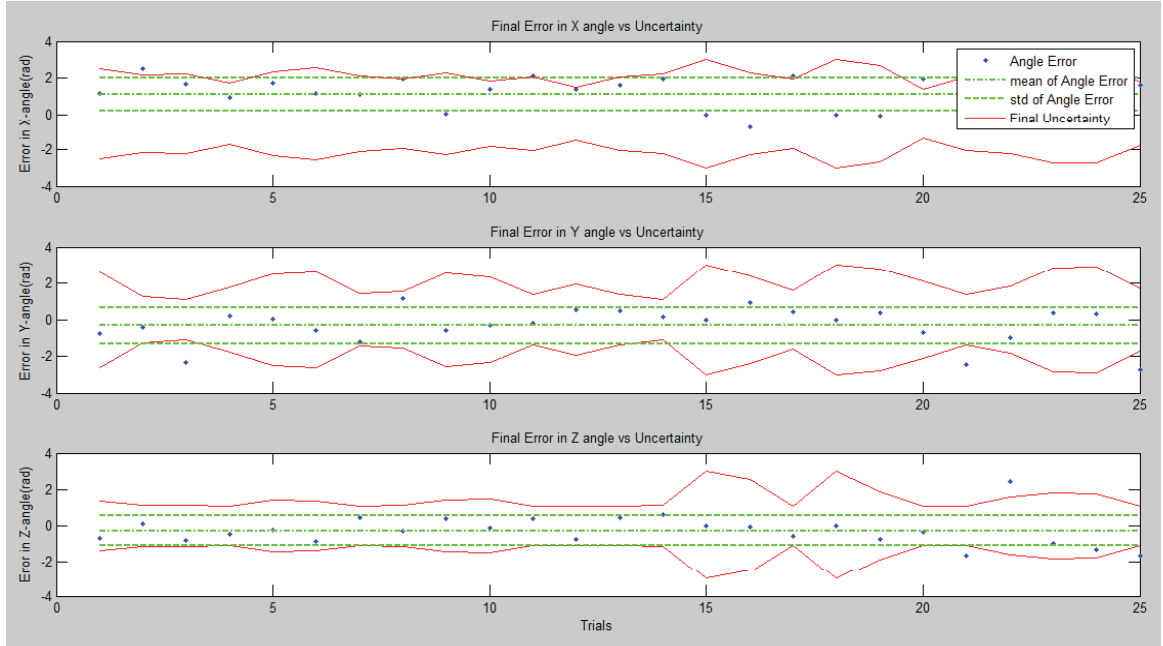| Geometry	Elevation	Azimuth	# of Meas.
1	25	35	2
2	35	60	3
3	30	45	3
4	60	60	3
5	35	35	2
6	20	35	3
7	45	65	3
8	60	70	3
9	70	40	2
10	60	30	2
11	40	50	3
12	65	40	3
13	40	55	3
14	30	55	3
15	25	55	2
16	65	65	1
17	55	65	3
18	15	30	0
19	25	70	2
20	70	40	3
21	40	35	3
22	65	15	2
23	20	70	1
24	15	70	1
25	55	35	3

Figure 5.27: Error in Final Attitude Estimates for the 25 DIRSIG Geometries. This graph shows how a set of correct initial conditions for attitude are pulled off by incorrect measurements of β and θ.

Figure 5.28 shows the same results as Figure 5.27, but for ten trials at each image. In comparison, Figure 5.29 shows the results of adding uncertainty in the surface parameters at each geometry. Ten tests were run on each image with a 10% random error in each surface parameter. This figure still shows clusters of erroneous final estimates similar to the ones in Figure 5.28. However, most of these errors in the x and y axes now have a small spread. It can be seen, however, that this small spread is not close to the spread of errors due to geometry. It can also be seen that these types of errors hardly affect the z axis esimatiton. This is due largely to the fact that most of the angle of polarization, or θ measurement, is applied to the z axis angle and the θ measurement is not affected by errors in surface parameters.

Finally, Figure 5.30 shows the results from using a corrupted initial estimate of camera orientation in the filter. The final spread of error is slightly worse than the ones shown in Figure 5.29, but falls mostly in line with the expected uncertainty from
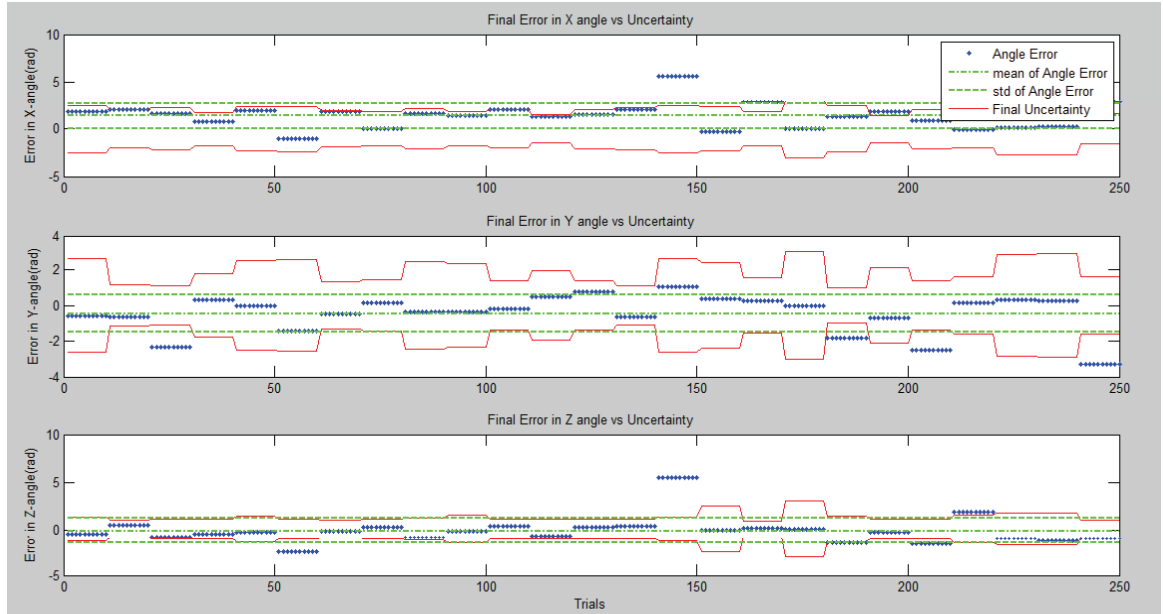
Figure 5.28:   Error as a Function of Geometry with Multiple Tests. This figure shows that with the same initial orientation and same measurements, the final errors are the same. This figure should be compared with Figure 5.29 which shows the final errors given random errors in surface parameters as well.
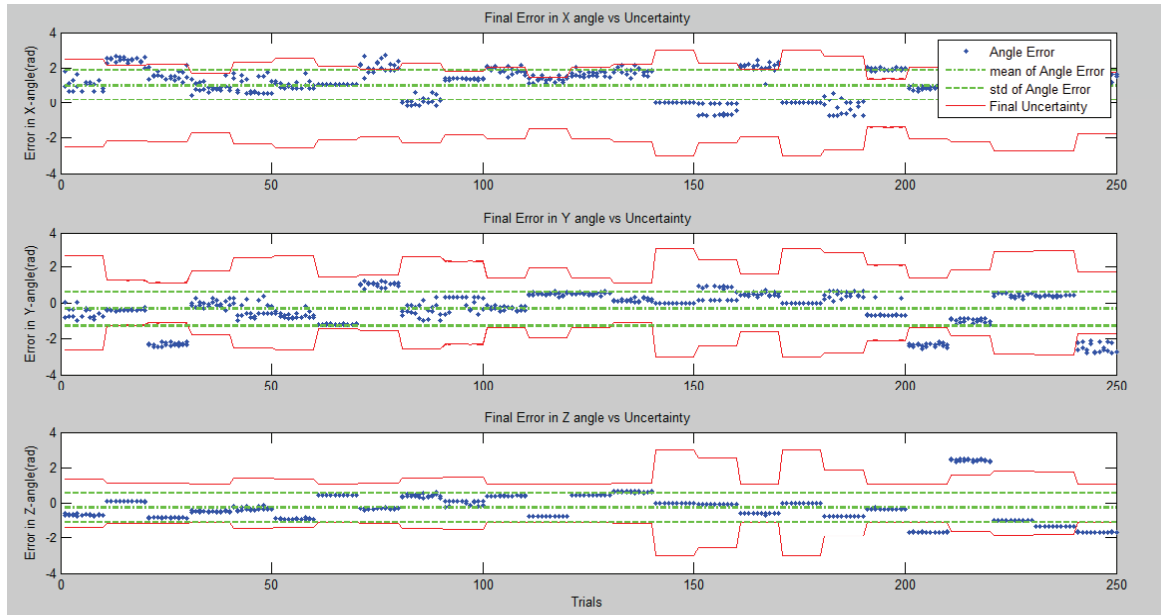
Figure 5.29:   Error as a Function of Geometry and a 10% Error in Shell Parameters. This figure shows that errors due to errors in surface parameter estimation are not as bad as errors due to bias in the estimation of the β angle.

133

Figure 5.30: Errors as a Function of Random Initial Error in Orientation, Bias in Estimation of β and a 10% Random Error in Shell Parameters. This shows final uncertainties close to the expected uncertainties from the Kalman Filter.

the filter. Table 5.9 shows the results for mean error in each angle, standard deviation in error from the Monte Carlo runs, and expected uncertainty from the filter.

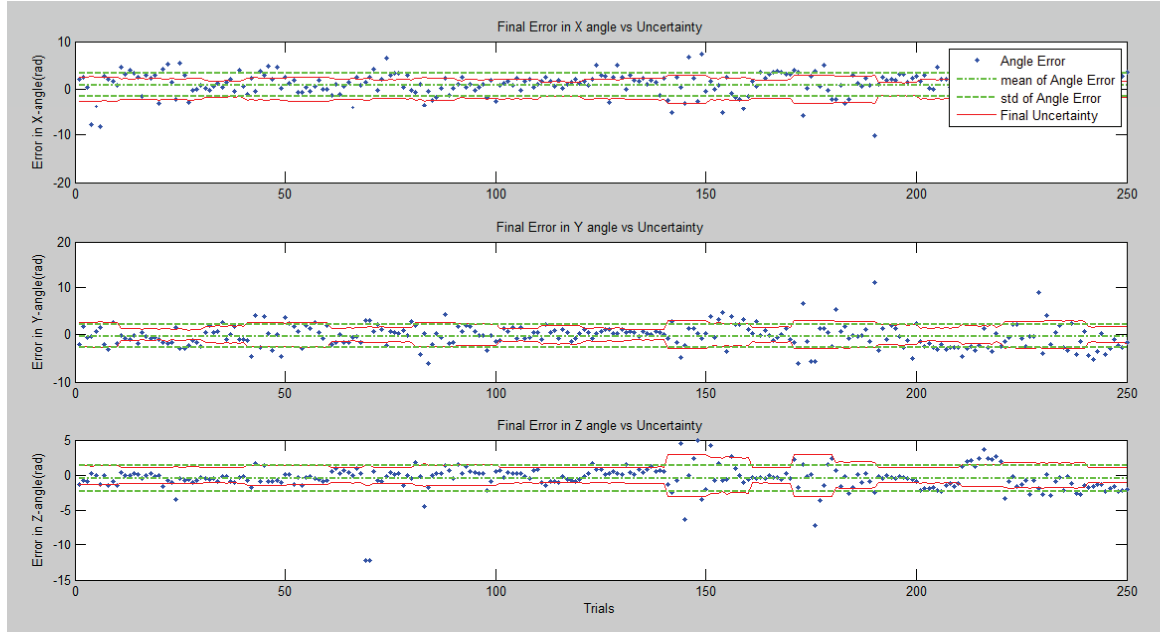These results show an average decrease in error of about 25% in the x and y axes and about 50% in the z axis. There is a bias in the x axis of almost 1°, but this might be mitigated by removing bias in the β measurement in future work.

Table 5.9: Final Error Results for the DIRSIG Simulations. These results show that there is a mean error in the x-axis of almost 1°, but that otherwise, the Monte Carlo errors and uncertainty from the Kalman filter match. These results show an improvement in error of about 25% in the x and y axes and about 50% in the z-axis.

	x-error (deg)	y-error (deg)	z-error (deg)
Mean Final Error (deg)	0.88	-0.30	-0.46
Std of Monte Carlo Runs (deg)	2.47	2.35	1.79
Kalman Filter Uncertainty (deg)	2.48	2.59	1.36

Table 5.10:   Table of Final Error Angles for the Physical System Test. This table shows a bias in the z-axis estimate of about $-5°$ due to a misalignment of the polarizer in the rotation stage. Otherwise the mean and standard deviations of these errors fall in line with the DIRSIG tests and the expected errors from the Kalman filter.

Test Run	x-error (deg)	y-error (deg)	z-error (deg)
1	0.80	1.42	-5.72
2	0.29	-1.20	-5.26
3	-2.76	3.06	-5.96
4	2.95	0.85	-4.53
5	-2.72	-0.56	-6.23

5.3.3 Physical System Test Results .   The limited physical system tests, described in Section 4.4.2.2, show similar results to the DIRSIG test. Table 5.10 shows the results of five test orientations completed with the physical system.

The first observation to be made is that the errors in the z axis appear to be much larger than those found in the DIRSIG tests. However, the spread of these errors is about the same. There is a bias in the angle of polarization measurement of almost $5°$, due to the coarse alignment of the polarizer in the rotation stage. Beyond the polarizer bias, it appears that the physical system test and the DIRSIG simulation are in line and that there is improvement in attitude estimation from degree of polarization measurements.

Ideally, more measurements would be made with proper random errors placed on all unknown components. However, given the time constraint and the nature of the simulation software, these types of tests are left to future work. Ideas on further research and expected results are covered in the next chapter.

VI. Conclusions and Future Work

This thesis presented a way to determine camera orientation from a single view using a passive polarimetric sensor. This chapter wraps up the thesis by discussing conclusions made from the results in Chapter V (Section 6.1). It goes on to discuss some thoughts on the limited testing conditions with the DIRSIG simulation software and the physical system, their implications in the results and how the full set of tools developed in Chapter III may be used to draw further conclusions (Section 6.2). It then presents ideas and expected impact from future work (Section 6.3).

6.1 Conclusions

Though the main focus of this thesis was the improvement in navigation solution using a passive polarimetric sensor, many conclusions were made along the way. It was shown that the interrelationship between geometry and surface parameters limits the full estimation of material and relative orientation. However, it was also shown that, with information about one of these, the other may be estimated well. These results were then used to develop a set of constraints on a material that would allow for the structure to be determined without knowing the full set of Shell parameters. This algorithm worked well under certain conditions but was improved upon when the full set of Shell parameters was known.

Given that a full set of Shell parameters could be determined from a catalog using the multiple hypothesis technique, and that the Kalman filter would provide an adequate geometry estimation for the multiple hypothesis test to work, a material was chosen from the catalog of known Shell materials. The chosen material was based on material parameters that would provide a good candidate material to test the Kalman filter algorithm.

In general, the results presented in Chapter V show that a this algorithms works well with materials and geometries with a fairly large degree of polarization, a broad spread in the degree of polarization, from a large σ parameter, and low diffuse and volumetric parameters. These parameters showed the most impact on degree of

polarization measurements when errors were present in either the estimate of the set of Shell parameters or estimated geometry.

6.2   Ideal Testing Conditions, Caveats and Impact on Results

A word of caution must be made about the results derived from the testing conditions used for the DIRSIG and physical system tests. It is admitted that these test conditions are limited. The nature of the complex equations and sheer number of parameters to test makes a full Monte Carlo test a hefty task.

Though the DIRSIG simulation software is easy to use for particular conditions, the nature of the software does not make it easy to test multiple random conditions. The physical system conditions are also not ideal. The limited availability of materials with known Shell parameters makes it difficult to test multiple materials. The balance required between the Vicon camera system and largely scattered lighting conditions has to be taken into account when setting up usable test conditions that require truth data. Also, non-perfect reflecting materials used to diffuse and scatter the sources when using the physical system can also be a hindrance.

Although the main focus of this thesis was aiding camera orientation estimates using passive polarization imaging, the way it was achieved was through a ground up research method. The results presented in Section 5.3 are correct for those particular conditions, but they can not be used to determine how well the Kalman filter algorithm would work under other conditions. However, the full set of tools developed throughout the research process can easily be used to determine expected results for another situation. This type of easy simulation was intentional and saves time in determining which situations are useful before undertaking the task of setting up a full DIRSIG or physical system test.

6.3   Future Work

Given additional time, the following endeavors would be of interest for continued work. There were several constraints used throughout the work that were required to make useful simplifications for specific tests. However, some of these constraints may be able to be relaxed in future work with the incorporation of other related research.

Using a camera that can capture multiple polarizer orientations and calculate degree and angle of polarization in a single snapshot would be useful in integration with other navigation aid sensors. A camera such as the ones found in [5] could be combined with inertial sensors to complete the Kalman filter cycle and determine update accuracy at a larger number of conditions.

Longer wavelengths of light may have benefits for particular scenarios. Wavelengths in which sources are predominantly thermal would mean that there should be a specular reflection under all geometries. These longer wavelengths are also more resilient to rough surfaces since surface roughness is a function of wavelength and surfaces becomes more specular at longer wavelengths. Also, according to the Shell parameter table found in [31], the diffuse and volumetric components of the model tend to become smaller at larger wavelengths. This would be ideal since it was shown in Section 5.1.1 that these two parameters have a large influence on degree of polarization measurements and that smaller values are more satisfactory.

These images are not only useful for degree and angle of polarization measurements. Other electro-optically aided navigation techniques can be used on the same images. Many of these techniques are useful in areas where the techniques presented in this thesis break down and vice versa. A few specific EO-aided navigation techniques that could be useful additions to this Kalman filter algorithm include feature matching techniques using epipolar constraints, which use areas of dense features, homographic constraints, which use flat areas, and vanishing point detection algorithms, which use lines in areas that may not have dense features and may not be flat.

Finally, the corresponding requirements between the Kalman filter and the multiple hypothesis testing algorithm present an opportunity for a symbiotic relationship. The requirement for the Kalaman filter to have a full set of Shell target parameters can be fulfilled by the multiple hypothesis algorithm. Likewise, the multiple hypothesis algorithm's need for an accurate relative geometry can be fulfilled by the Kalman filter. The two methods may then be able to be combined in a simultaneous localization and mapping technique. This method could be further improved through implementation of a correspondence algorithm, which would allow for multiple measurements of the same surface from different geometries.

6.4 Summary

The results of this thesis show that a simple adaptation to a readily available camera system can be used as an additional measurement for the purposes of attitude estimation. This information is only a building block in between existing algorithms and future integration techniques. However, these tools showed great promise and demonstrate the need for further exploration into EO-aided navigation techniques using polarization measurements.

Appendix A. Material Parameters

This appendix presents a list of the materials and their associated Shell target parameters [31].

Table A.1:   Material Properties. This table shows the Shell target model parameters for each of the materials used throughout this thesis. A description of each parameter can be found in Chapter II.

Material	(n)	(k)	(B)	(σ)	(τ)	(Ω)	(ρ_D)	(ρ_V)
White Paint	1.515	.112	.022	.008	0.134	1.459	0.364	-5.01×10^{-1}
Concrete	1.498	0.4071	0.2644	0.8574	55.36	0.0606	2.29×10^{-2}	2.25×10^{-2}
Tan Paint	1.43	0.3573	0.1093	0.8029	58.79	52.39	0.1114	2.31×10^{-2}
Green Paint	1.39	0.3371	0.1048	0.4563	18.54	36.57	6.914×10^{-3}	1.552×10^{-3}
Aluminum	5.92	0.3045	0.129	0.0018	0.2145	4.639	5.16×10^{-3}	3.466×10^{-3}
Flat Black Paint	1.405	0.2289	0.0056	0.3331	1.717	119.3	-1.762×10^{-4}	5.427×10^{-4}
Glossy Black Paint	1.4	0.4	1.3	0.05	5	10	1.1×10^{-5}	10^{-7}

Bibliography

1. URL www.newport.com.

2. URL www.vicon.com.

3. *Unscented Filtering and Nonlinear Estimation*, volume 92 of *Proceedings of the IEEE*, 2004.

4. et. al., Shih-Schon Lin. "Polarization Enhanced Visual Surveillance Techniques". *Proceedings of the 2004 IEEE International Converence on Networking, Sensing and Control*. 2004.

5. Applied Technology Associates. *The Polarized Infrared Advanced Tactical Experiment*. Technical report, Air Force Research Laboratory VSSS, 2006.

6. Bouguet, Jean-Yves. "Camera Calibration Toolbox for MATLAB". URL http://www.vision.caltech.edu/bouguetj/calib_doc.

7. Brown, Alison and Dan Sullivan. *Inertial Navigation Electro-Optical Aiding During GPS Dropouts*. Technical report, NAVSYS Corporation, 2002.

8. Caughlan, James M. "Manhattan World: Compass Direction from a Single Iamge by Baysian Inference". *Computer Vision 1999: The proceedings of the Seventh IEEE International Conference on*. IEEE International Conference on Computer Vision, 1999.

9. Chi, Jinkui and Kaichun Zhao. "Design of a Novel Polarization Sensor for Navigation". *Proceedings of the 2007 IEEE International Converence on Mechatronics and Automation*. 2007.

10. d'Angelo, Pablo and Chrisitan Wohler. *3D Surface Reconstruction Based on Combined Analysis of Reflectance and Polarisation Properties*. Technical report, Daimler Chrysler Research and Technology, Machine Perception, 2005.

11. d'Angelo, Pablo and Chrisitan Wohler. "3D Surface Reconstruction by Combination of Photopolarimetry and Depth from Defocus". *In Pattern Recognition, Proceedings of the 27th DAGM Symposium, LNCS 3663*. 2005.

12. d'Angelo, Pablo and Chrisitan Wohler. *Image-based 3D surface reconstruction by combination of sparse depth data with shape from shading and polarisation*. Technical report, Daimler Chrysler Group Research, Machine Perception, 2005.

13. Egan, Walter. *Photometry and Polarization in Remote Sensing*. Elsevier Science Publishing Co, Inc., 1985.

14. Gartley, Michael G. *Polarimetric Modeling of Remotely Sensed Scenes in the Thermal Infrared*. Ph.D. thesis, Binghamton University, 1995.

15. Gonclaves, Debora and Eugene Irene. "Fundamentals of Applications of Spectroscopic Ellipsomentry".

16. Hartley, Richard and Andrew Zisserman. *Multiple View Geometry in Computer Vision.* Cambridge University Press, 2nd edition, 2003.

17. Hyde, Mylo. *Determining the Index of Refraction of an Unknown Objects Using Passive Polarimetric Imagery Degraded by Atmospheric Turbulence.* Ph.D. thesis, Air Force Institute of Technology, 2010.

18. Koshikawa. "A Polarimetric Approach to Shape Understanding Of Glossy Objects". *International Joint Conference on Artificial Intelegence*, 1979:493–495, 1979.

19. Lourakis, Manolis and Antonis Argyros. "Efficient, causal camera tracking in unprepared environments". *Computer Vision and Image Understanding*, 99:259–290, 2005.

20. Lowe, David. "Object Recognition from local scale-invarian fatures." *Proceedings of the International Conference on Computer Vision.* 1999.

21. Lowe, David. "Distinctive Image Features from Scale-Invariant Keypoints". *International Journal of Computer Vision*, 60:91–110, 2004.

22. Marquardt, Donald. "An Algorithm for Least-Squares Estimation of Nonlinear Parameters". *SIAM Journal on Applied Mathamtics*, 11:431–441, 1963.

23. Maxwell, J.R., J. Beard, S. Weiner, D. Ladd, and S. Ladd. *Bidirectional Reflectance Model Validation and Utilization.* Technical Report AFAL-TR-73-303, Environmental Research Institute of Michigan, October 1973.

24. Maybeck, Peter. *Stochastic Models, Estimation, and Control Volume I.* Academic Press - Mathematics in Science and Engineering, 1979.

25. van der Mere, Rudolph, Arnauld Doucet, Nando de Freitas, and Eric Wan. *The Unscented Particle Filter.* Technical report, Cambridge University Engineering Department, 2000.

26. Meyers, Jason. *Modeling Polarimetric Imaging using DIRSIG.* Ph.D. thesis, Michigan Technological University, 1993.

27. Michael J. Veth, Maj. *Fusion of Imaging and Inertial Sensors for Navigation.* Ph.D. thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, September 2006. AFIT/DS/ENG/06-09.

28. Priest, Rishard and Thomas Germer (editors). *Polarimetric BRDF in the Microfacet Model: Theory and Measurements*, volume 1. Proceedings of the 2000 meeting of the Military Sensing Symposia Specialty Group on Passive Sensors, August 2000.

29. Proceedings of SPIE, the International Society for Optical Engineering. *Surface Orientation from Polarization Images*, volume 850, 1987.

30. Rochester Institute of Technology. *DIRSIG User's Manual.* Http://dirsig.org/docs/maual-2006-11.pdf.

31. Shell, James R. *Polarimetric Remote Sensing in the Visible to Near Infrared.* Ph.D. thesis, Rochester Institute of Technology, 2005.

32. Titterton, David H. and John L. Weston. *Strapdown Inertial Navigation Technology.* The Institution of Eletrical Engineers, Reston, Virgina, 2004.

33. Torrance, K.E. and E.M. Sparrow. "Theory of Off-Specular Reflection From Roughened Surfaces". *Journal of the Optical Society of America*, 57:1105–1114, 1967.

34. Various. *Physics-Based Vision: Principles and Practice, Volume 1: Radiometry.* Bartlett Publishers Inc., 1992.

35. Various. *Physics-Based Vision: Principles and Practice, Volume 3: Shape Recovery.* Bartlett Publishers Inc., 1992.

36. Wolff, Lawrence B. "Polarization-Based Material Classification from Specular Reflection". *Physics-Based Vision: Principles and Practice, Volume 1: Radiometry.* Bartlett Publishers Inc., 1988.

37. Wolff, Lawrence B. (editor). *Surface Orientation from Two Camera Stereo with Polarizers*, volume 1194. SPIE, 1990.

38. Wolff, Lawrence B. and Terrance Boult. "Constraining Object Features Using a Polarization Reflectance Model". *Physics-Based Vision: Principles and Practice, Volume 1: Radiometry.* Bartlett Publishers Inc., 1990.

Index

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704–0188

1. REPORT DATE *(DD–MM–YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From — To)*
24–03–2011	Master's Thesis	Sept 2009 — Mar 2011

4. TITLE AND SUBTITLE

Polarimetric Enhancements To Electro-Optical Aided Navigation Techniques

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Jeremiah Johnson, Capt, USAF

5d. PROJECT NUMBER

None

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT/GE/ENG/11-19

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Office of Scientific Research
875 N Randolph St
Ste 325 Rm 3112
Arlington VA 22203
jon.sjogren@afosr.af.mil DSN: 426-6564

10. SPONSOR/MONITOR'S ACRONYM(S)

AFOSR/RSE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Navigation in indoor and urban environments by small unmanned systems is a topic of interest for the Air Force. The Advanced Navigation Technology Center at the Air Force Institute of Technology is continually looking for novel approaches to navigation in GPS deprived environments. Inertial sensors have been coupled with image aided concepts, such as feature tracking, with good results. However, feature density in areas with large, flat, smooth surfaces tends to be low. Polarimetric sensors have been used for surface reconstruction, surface characterization and outdoor navigation. This thesis combines aspects of some of these algorithms along with a realistic, micro-facet polarimetric model and a Kalman filter approach to determine surface structure and platform orientation in an indoor environment. A series of graphical user interfaces were developed to estimate surface material parameters. A more complex simulation software package was used to estimate camera attitude. A physical polarimeter was also designed and built to test the algorithms in a realistic environment. An improvement in attitude estimation of up to 50% was demonstrated.

15. SUBJECT TERMS

Polarimetry, Navigation, Electro-Optically Aided

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Michael J. Veth, Lt Col, USAF (ENG)
U	U	U	UU	161	**19b. TELEPHONE NUMBER** *(include area code)* (850) 882–4667; michael.veth@eglin.af.mil

Standard Form 298 (Rev. 8–98)
Prescribed by ANSI Std. Z39.18